Development of a Prototype Domain-Specific Language for Monitor and Control Systems

Matthew Bennett, Richard Borgen, Klaus Havelund, Michel Ingham, David Wagner
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
818-393-6426
{matthew.b.bennett, richard.l.borgen, klaus.havelund, michel.d.ingham, david.a.wagner}@jpl.nasa.gov

Abstract—This paper describes the Domain-Specific Language (DSL) prototype developed for the NASA Constellation Launch Control System (LCS) project. A key element of the LCS architecture, the DSL prototype is a specialized monitor and control language composed of constructs for specifying and programming test, checkout, and launch processing applications for flight and ground systems. The principal objectives of the prototyping activity were to perform a proof-of-concept of an approach to ultimately lower the lifecycle costs of application software for the LCS, and to explore mitigations for a number of development risks perceived by the project. The language has been implemented as a library that extends the Python scripting language, and validated in a successful demonstration of capability required for Constellation.[12]

TABLE OF CONTENTS

1. INTRODUCTION

In 2004, the White House announced a new Vision for Space Exploration [1], calling for a campaign of ambitious missions involving human and robotic exploration of the Moon, Mars and beyond. In response to this call, NASA established the Constellation program [2], which is planning to replace the aging space shuttle fleet with new vehicles capable of transporting crew and cargo to the International Space Station, then on to Moon and, ultimately, Mars.

Constellation represents the first major human space flight program start in many years. Budget constraints will require the program to make gradual progress towards its far-reaching goals over an extended schedule, placing unprecedented requirements on the evolvability of the supporting infrastructure and challenging NASA to significantly reduce the operations and sustaining costs associated with its missions. Consequently, the program's success will require significant upgrades to the ground-based infrastructure needed to support the assembly, test, and operations of these new vehicles.

One such system is the Launch Control System (LCS) at the Kennedy Space Center (KSC). In addition to coordinating and controlling the launch sequence, this system will be used to test the spacecraft, launch vehicles, and possibly their component subsystems as they are delivered to the space center and assembled for launch, to control various pieces of ground support equipment used in those operations, and to ensure the safety of all of these operations. Within the LCS, software applications will command distributed hardware end items, including the vehicles and associated ground support equipment, and monitor the telemetry produced by these end items.

Under the auspices of the Constellation Ground Operations project, engineers at KSC, in collaboration with personnel from JPL and other NASA Centers, conducted an LCS proof-of-concept activity over a roughly one-year period from late 2006 to late 2007. This proof-of-concept effort was intended to explore ways to mitigate a number of development risks perceived by the LCS project. Key among those risks were projections of the costs associated with the development and maintenance of software applications used to conduct launch processing operations. Experience with the current Space Shuttle test and checkout system suggests that a significant part of the ongoing operations costs are related to the development and maintenance of monitor and control applications. Particularly expensive is the extensive process by which system engineers express requirements for test procedures

1

[1] 1-4244-1488-1/08/$25.00 ©2008 IEEE.
[2] IEEEAC paper # 1576, Version 2, Updated October 23, 2007

in prose, software developers translate these requirements into code, and then both sets of experts are engaged in verification of the resulting application's correctness. The new system will have to significantly reduce these costs while assuring a consistent high level of safety and security.

One element of the LCS proof-of-concept activity was an investigation of the potential benefits of using a Domain-Specific Language (DSL) that systems engineers would be able to use to write *executable specifications* of monitor and control applications (i.e., capture detailed requirements in a form that would either be directly executable or automatically translatable to software implementation). The principal objective of this DSL prototype was to demonstrate an approach that will ultimately lower the lifecycle costs of monitor and control applications for launch processing, test and checkout.

This paper presents the results from the LCS DSL prototyping task. Section 2 introduces the problem domain and provides a brief overview of the envisioned LCS architecture. The paper goes on to describe our process for design and development of the LCS DSL, including a survey and assessment of existing DSLs against the LCS requirements and evaluation criteria (Section 3), a description of the language design (Section 4), an overview of the prototype implementation (Section 5), and a discussion of how the language and prototype applications were verified and validated (Section 6). Finally, the paper concludes with a discussion of observations and issues raised during the DSL prototyping effort.

2. PROBLEM DOMAIN & SYSTEM ARCHITECTURE

The LCS problem domain has a number of defining characteristics that pose important challenges and constraints on the system architecture. Given that human lives are at stake, there are especially high demands for stability, security, verifiability and fault tolerance. With the routine handling of hazardous materials, it is essential to physically distribute the system to separate human operators from potentially catastrophic events. In addition, given the large past investments in physical infrastructure, significant interfaces to legacy systems are unavoidable. Finally, the Constellation program represents a massive multi-center, multi-decade effort, so the new LCS will also be a very long-lived system, making total life-cycle costs a major concern. Such overarching issues drive the qualities required of the LCS architecture.

Prior to the start of the proof-of-concept activity, the project evaluated several architecture options. The existing architectures for the Space Shuttle, the International Space Station, and the EELV systems were not ultimately selected, due to a number of distinguishing requirements for the Constellation program and lessons learned from the experiences of developing and maintaining these systems. Nonetheless, various aspects of these architectures served as a source for the significant subset of requirements that the Constellation LCS will share with these previous systems. For example, the scripting language used in the current Space Shuttle command and control system, called the Ground Operations Aerospace Language (GOAL) [3], was used as a basis for comparison during the DSL survey and assessment, and as a source of inspiration for the specification of the LCS DSL.

At the end of the architectural evaluation process, the Constellation Ground Operations project opted to develop a "standards-based architecture", that is, an architecture that calls for extensive use of industry standards and standard software interfaces to achieve interoperability, adaptability, portability and lower development costs. More broadly, the standards-based architecture calls out the following architectural principles:

- Distribution of control responsibilities, especially to separate supervisory vs. real-time control;
- Incorporation of standard industrial controls and processes;
- An adaptive strategy that permits variations or extensions with commercial or legacy components;
- Decoupling of software components with communication middleware;
- Decoupling of software components by use of layering.

The purpose of the LCS system is to support testing and checkout functions as a spacecraft or other vehicle hardware is processed from arrival at KSC through launch. These functions include integrated tests to assure launch readiness. Major elements of the system would be distributed geographically across the KSC launch site, connected by institutional networks, as illustrated in Figure 1. These elements are divided into five major categories:

- ***End Items*** include the vehicle, launch pad, their subsystems, and other test articles and support equipment that are the targets of control. Different test and checkout applications may interact with different end items.
- ***Gateways*** provide hardware and software interfaces between the LCS system network and the end items, including the formats needed to communicate with particular end items.
- ***Applications*** are the control programs that direct and coordinate (or supervise) control of the end items. Control includes the evaluation and maintenance of

system health and safety, and performing checkout procedures needed to verify proper assembly and configuration.

- ***Displays*** provide operator feedback and a mechanism for operators to interact with the applications and, indirectly, with the end items.

- ***Middleware*** provides a data distribution fabric that isolates applications from the details of network protocols, data formats, and other programming details.

In part, this distributed architecture addresses a physical system constraint: because checkout and launch operations can be hazardous, the launch facility maintains significant distance between areas where those operations are conducted, and areas where the control computers and operators supervise them. Leveraging commodity (and in some cases, existing) networking capabilities to support the bulk of the connectivity would significantly reduce the hardware costs compared to any custom-built solutions. However, many of the target end items would not be able to directly support standard network connections or protocols. To bridge the gap between the LCS and various end items, the architecture provides gateways whose job it is to translate protocols between the LCS and middleware standards, and the diverse media and protocols needed to communicate with the end items. External protocols might include various supervisory control and data acquisition (SCADA) interactions with embedded controllers or industrial controllers, space telecommunications protocols such as those specified by the Consultative Committee for Space Data Systems (CCSDS), and various system bus protocols such as MIL-STD-1553.

Supervisory control would be conducted from test-specific software applications running in application servers in a secure control facility. A key architectural principle is the division of control responsibilities between supervisory control applications and the end items themselves, based on real-time constraints. Because of the distributed nature of the system, and the technology used to support standard network communications and processing, there are limitations to the ability of these networks to deliver data with the kind of deterministic, low-latency timing needed to support real-time closed loop control. Thus, the architecture defines supervisory control to exclude any closed loop control that would require reactions within time constraints shorter than a given threshold (as of this writing, the specific threshold has yet to be determined).

The LCS architecture draws a separation between application servers and control room computers used to host displays and perform user interface functions. This division is intended to enable supervisory control applications to interface with operators indirectly, allowing, for example, one application to interface with multiple operators (operator consoles, or operator stations in the control room), and to allow the user interface functions to be redirected to other displays in the event of a display hardware failure without affecting the continuity of the control application.

This is just one example of many techniques employed to achieve reliability and safety. The architecture calls for various protections against operator errors, including automation of sequencing, automated rule enforcement, appropriate design of human interfaces and multiple levels

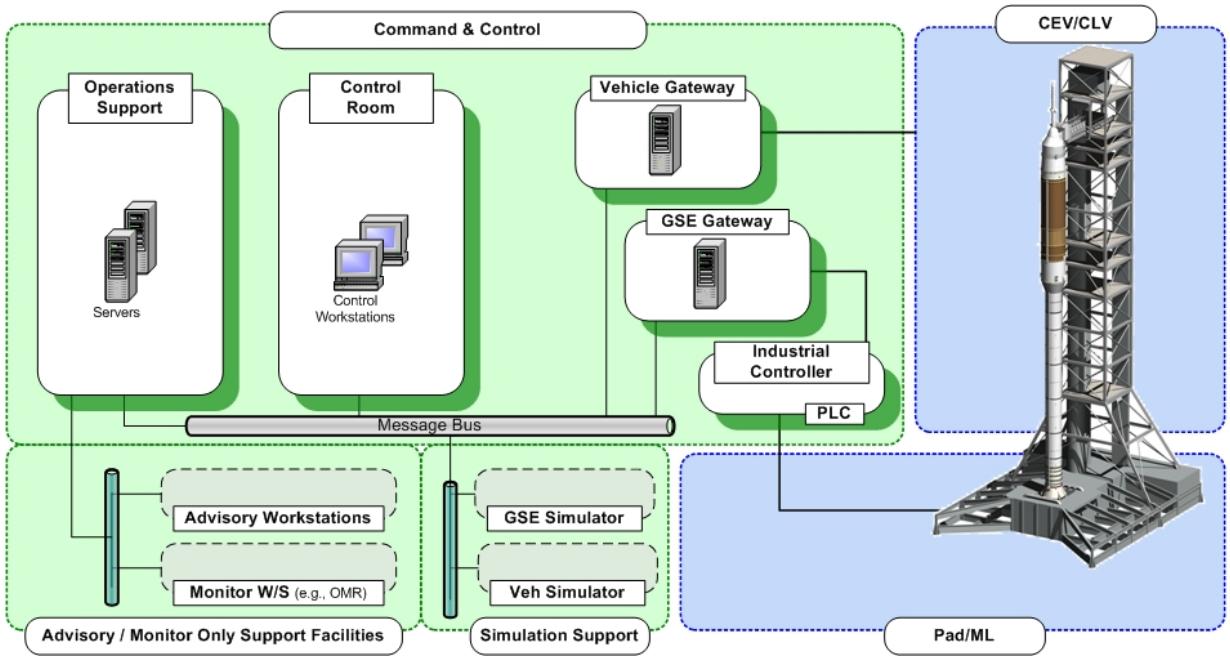

Figure 1. The Constellation LCS distributed architecture.

of automatic safing. Other measures are aimed at achieving a high level of safety for software and equipment, including security, reliable components, redundancy, and high availability techniques such as fault recovery.

A fundamental principle underlying the LCS architecture is the decoupling of software components through layering, as shown in Figure 2. Although IP networking hardware and protocols are ubiquitous, they provide a rather primitive level of service. A message-oriented middleware (MOM) layer is used to provide higher-level application services such as data distribution using a publish-subscribe mechanism, and an abstraction of the logic needed to govern the qualities of service. Abstracting the management of data delivery allows applications to control the trade-offs between data latency, continuity, and volume, without having to get involved in the details of how these services are rendered on underlying hardware and network protocols.

A key characteristic of the publish-subscribe delivery pattern is that it decouples distributed applications from having to know all of the details about the system topology such as which applications need to receive information that a given application creates. The application can publish the message on a defined *topic*. Other applications needing to use that data can subscribe to the topic, and the middleware will manage the distribution and delivery of the data from the publisher to all subscribers according to the qualities of service defined for the topic.

To further isolate applications from the platform and middleware details, the architecture calls for a common services software framework layer. This layer defines the message topics and default qualities of service that the system will support, and adds some topic-specific services that support system-level interactions between applications.

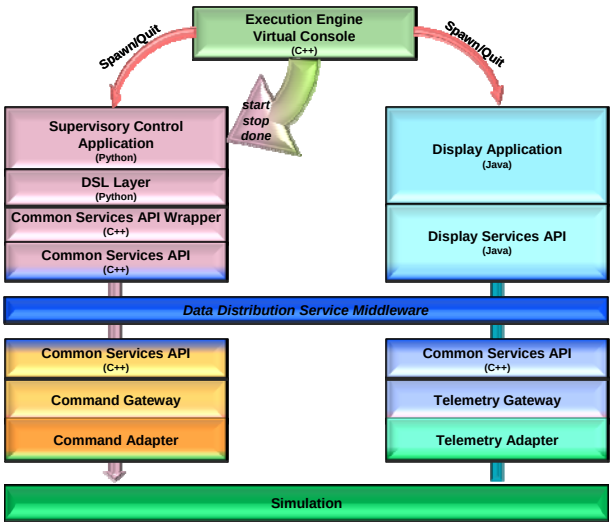

Figure 2. Layered view of the LCS architecture.

Additional frameworks are defined on top of this layer to provide support for specific kinds of applications (e.g., display applications, gateway applications, monitor and control applications).

The proof-of-concept demonstration required the selection of concrete technologies and products to instantiate a representative prototype of the target architecture. For the MOM layer, the project used a commercial implementation of the Data Distribution Service (DDS), a standard defined by the Object Management Group (OMG) [4]. Another commercial product was used to provide gateway services supporting telemetry and commanding. The remainder of this paper focuses on the design of the Supervisory Control, DSL, and Common Services layers of the architecture.

3. LANGUAGE SURVEY

One of the early investigations performed as part of the LCS proof-of-concept activity was a study of various programming languages appropriate for use in monitor and control applications. Following a broad survey of available GOTS/COTS DSL and general-purpose language solutions, a small set of candidate languages were selected for more detailed investigation, including six DSLs (with varying degrees of domain-specificity; specific references are omitted here) and three general-purpose programming languages (Python, C++, and Java). Each candidate language was assessed against a set of evaluation criteria. Functional requirements, life cycle costs, and various quality metrics were considered in the assessment process, which included prototype development and analysis for each candidate language. Prototyping of an example test scenario provided the hands-on experience needed to support the analysis results. This survey and assessment was intended to feed a "make versus buy" decision for the monitor and control language to be used in LCS application software development. This section of the paper describes the requirements and evaluation criteria used in the assessment and summarizes the results from this investigation.

Requirements

The requirements were organized into categories and were used to determine if a particular language could meet the needs of LCS:

- ***General*** included requirements for applications to be interoperable with C and C++ system software, be able to perform interruptible and non-interruptible delays, and invoke other applications that execute in parallel or in series.

- ***Commanding*** included requirements for applications to be able to command ground support equipment and flight end items.

- ***Event Handling and Conditional Execution*** included requirements for applications to be able to define logical, numeric, and temporal conditions, and to perform conditional branching, immediate verification of conditions, verification of conditions within a time period, and continuous verification of conditions. There were requirements to be able to define events based on measurement value and attribute changes, time changes, external event notifications, and user inputs. In addition, there were requirements to subscribe and unsubscribe to events, and to respond to events with behaviors that include sending text messages to users, computing derived measurements, sending commands, invoking other applications, and initiating safing programs. Finally, there were requirements to enable and disable events.

- ***I/O Interface*** included requirements for applications to be able to prompt users for input, to be notified of user responses to prompts, and to publish or display a textual message.

- ***Communication*** included requirements for applications to be able to communicate with one another, and be able to subscribe to only specified message topics.

- ***Data Handling*** included requirements for applications to be able to read the last known value, status, and attributes for a measurement, and to update a measurement's value and attributes. In addition, there were requirements to compare measurement data against criteria using equal, not equal, greater, greater than or equal, lesser, lesser than or equal, and range comparators. Finally, there were requirements to compare measurement data against specified criteria within a specified time period, and to compute criteria that are numeric and Boolean functions of other measurement data.

Evaluation Criteria

The criteria were organized into categories and were used to qualitatively evaluate how well a particular language meets the needs of LCS:

- ***Usability*** included criteria for simplicity, readability, ease of development, configurability of the code generator for compiled languages, modularity, ease of integration, ease of installation and configuration of an integrated development environment.

- ***Reliability*** included criteria for verifiability, validatability, diagnosability, and ease of troubleshooting.

- ***Maturity*** included criteria for commercial availability, pre-existing NASA and DoD usage, size of existing major development efforts, availability of implementations, and availability of support from commercial and user community sources.

- ***Performance*** included criteria for resource usage, responsiveness for reactive control, and scalability.

- ***Versatility*** included criteria for applicability to embedded applications as well as non-embedded applications.

- ***Extensibility*** included criteria for the ability to add new capabilities such as new APIs and libraries, extensions to the language syntax, and new modules and/or parameterized procedures.

- ***Portability*** included criteria for loose coupling with a particular toolset or execution environment, and for level of effort to adapt the language to the LCS computing platform.

- ***Lifecycle costs*** included licensing, infrastructure setup and maintenance, and perceived costs for application development (including the use of tools).

Survey Results and Conclusions

None of the evaluated languages seemed to be a "100% perfect" fit for the specific requirements of the Constellation LCS (for example, as compared to the GOAL DSL in the Space Shuttle ground operations system today [3]). That is to say, none of them ranked "high" against all of the established evaluation criteria. However, any of the surveyed domain-specific or general-purpose languages could be used to satisfy the functional requirements, though in some cases it might require significant effort, and result in a non-optimal system from the standpoint of the non-functional evaluation criteria. Although the domain-specific languages offer syntax that is suited to monitor and control applications, they also tend to have more limited development and/or execution environments and tool support. General-purpose languages offer significant flexibility in terms of implementing required functionality and impressive portability and maturity characteristics, but suffer primarily in terms of readability, verifiability and validatability, considering the targeted user base of systems engineers.

Considering the risk-reduction objectives of the LCS proof-of-concept activity, and the uncertainty and fluidity still associated with the Constellation Program and LCS Project requirements, a two-pronged strategy for the prototyping phase of the LCS proof-of-concept was recommended, in which both a COTS solution and a "home-grown" solution would be explored. This would enable the LCS proof-of-concept team to mitigate the risk associated with premature down-selection of a single application specification language that may not end up satisfying all requirements.

Based on the team's technical assessment, one particular DSL among the COTS options was identified as the strongest contender due to its human space flight qualification heritage, and generally favorable usability and portability characteristics. However, the focus here is on the design of the "home-grown" DSL, so discussion of the COTS solution is out of scope for this paper.

Among the general-purpose programming languages considered, Python [5] was seen as the most suitable for use as a monitor and control language, based on the evaluation criteria established for the LCS Project. Python was selected to be the base language for the home-grown prototype DSL solution, because it presents the smallest "semantic gap" and shallowest learning curve to a systems engineer user, among the general-purpose programming languages considered. This being said, a couple of areas of particular concern were identified regarding Python, which were examined during the prototyping activity to determine if either of these present a showstopper obstacle for the ultimate adoption of a Python-based DSL. The two areas were Python's lack of static type checking, and the uncertainty associated with the ability of Python to satisfy performance requirements.

Internal vs. External DSL Approach

In implementing a home-grown DSL solution based on the Python programming language, two competing strategies were considered:

- ***Internal DSL Approach:*** that is, extending the base language with domain-specific constructs. In this approach the DSL is Python extended with a LCS-specific Python library. Applications would be written in Python and the execution engine would be that of Python.

- ***External DSL Approach:*** that is, translating a custom DSL into the base language. In this approach the DSL would be distinct from Python, adopting a syntax appropriate to the target domain. The DSL would be translated to Python, and the execution engine would be that of Python, as in the internal DSL solution.

For the LCS proof-of-concept, the team decided to implement an internal DSL. Given the short timeframe for the proof-of-concept activity, it was determined that this approach presented the least schedule risk; the existing Python development environments and tools could be used to facilitate the quick production of a prototype.

Beyond the LCS proof-of-concept, the team issued a recommendation to consider the development of an external DSL and toolset.[3] There are significant long-term benefits for having a language customized for LCS applications. In particular, an external DSL has the greatest potential for reducing lifecycle costs. Applications should be more readable/reviewable, verifiable, easier to write by systems engineers, and ultimately more reliable. However, the external DSL option would require significantly more work than the internal option; namely, designing the language, building a translator, constructing a development environment, and building tools such as a debugger, static analyzer(s), and a test harness.

[3] As discussed in the next section, the LCS team at KSC subsequently developed a prototype Tabular DSL, which is a form of external DSL built on top of the LCS (internal) Python DSL. This Tabular DSL has since been integrated with the rest of the LCS proof-of-concept system.

4. LANGUAGE DESIGN

The DSL is intended to facilitate writing control applications that remotely monitor and control the state of a set of end items. From the point of view of a control application, the state can be understood as a mapping from measurement variable names (each associated with a sensor in an end item) to value assignments. In the LCS architecture, this mapping is provided through the MOM software sitting between the control application and the end item.

The prototype DSL has been implemented as a Python library consisting of a set of functions (and a few classes). The design is inspired by traditional procedural programming, rather than by object oriented programming, in spite of the fact that Python is an object-oriented programming language. That is, the emphasis is on functions rather than on classes and objects. This design was influenced by the goal of defining a domain-specific programming language, where each function in the library represents a language construct. In addition to these domain-specific constructs, all of Python's programming constructs are available to the programmer. In the following discussion the key language constructs are presented:

Measurements

Measurements are views into the state of end items that are being monitored; they represent data samples collected from sensors in the end items. These samples are then drawn into the LCS system network via a gateway, and distributed as measurement messages on the middleware message bus. A measurement service buffers one or more of the most recent samples from each sensor so that these messages can be retrieved on demand by the application. The measurement service design is discussed in more detail in Section 5 of this paper.

A measurement object is an instance of the following class, for which only some of the methods are shown:

```
class Measurement:
  def __init__(self, id, value)
  def getId(self)
  def getValue(self)
  def getTime(self)
```

```
def __lt__(self, other)
def __eq__(self, other)
def __add__(self, other)
def __sub__(self, other)
…
```

A measurement object contains a time tag (defined using the Posix Time System), a numerical identifier, and a value. The time tag is automatically inserted in the object upon creation. The class is able to internally represent different value types, and the class also contains methods for examining the type and for extracting values depending on their type. The class defines a set of mathematical relational methods for comparing values (`__lt__` , `__eq__` , …), corresponding to the relational operators `<` , `==`, etc. The method `__lt__` for example has the definition:

```
def __lt__(self, other):
  return self.value < other.value
```

The methods are named in such a way that they overload the built-in relational symbols. For example, given two measurements m_1 and m_2, they can be compared using traditional syntax:

```
if m1 < m2: …
```

Measurements can be accessed by the name represented as a string, by calling the function **`getByName`**`(name)`, for example: **`getByName`**`("pressure")`. An alternative approach is offered through a specially engineered object named **`read`**: the same measurement can be accessed by the term **`read`**`.pressure`, without having to indicate quotes[4].

Monitoring Functions

The DSL monitoring functions offer capabilities for testing the values of named measurements as a function of time. A monitoring function is characterized along five dimensions:

Condition: the condition on a subset of the monitored measurements. For example: "pressure >= 300".

Timing: indicating when the condition should be verified to be true. There are three possibilities:
(i) now;
(ii) continuously during a specified time period, given as an extra parameter; and
(iii) eventually within a specified time period, given as an extra parameter.

Reaction: a reaction to be executed in case the property gets violated.

Repetition: a Boolean indicating whether the verification should continue even if the property gets violated.

Blocking: a verification can be specified to be *blocking* in which case the calling application will wait until the verification has been performed. Alternatively, the verification can be *non-blocking*, where the checking is "spawned" to the background while the calling application continues to execute. In a non-blocking case where a reaction is to be executed upon violation of the condition, there is a further choice between letting this reaction execute in parallel with the calling application or letting it interrupt the calling application.

Some of these monitoring functions are described in the following. The construct descriptions use the following symbols: `C` stands for a condition to be verified and `R` stands for a reaction to be executed in case a condition gets violated. Both `C` and `R` are assumed to be parameter-less functions. `D` stands for a duration, expressed in seconds. `S` stands for a string – generally a name associated with the verification operation. `N` stands for a natural number. Finally, `F` stands for a Boolean flag indicating whether verification should be repeated in case of property violations. Arguments in square brackets `[..]` denote optional arguments (note that this is not Python syntax). In Python, such optional arguments can be given default values, but these are not indicated here.

`verify``(C,[R],[S])`

This blocking construct examines whether the specified condition `C` is true **now**. The reaction `R` indicates the response that should be executed if the condition is not true. The default reaction is a dialog that is initiated with the user, giving the choice between iterating (redoing the verification), returning from the **`verify`** function, or aborting the control application entirely. The name `S` passed to the function helps identify the call within an execution log in case a property gets violated. The function returns true or false depending on whether the property succeeds. An example of the use of this function is provided below, where the condition and the reaction are first defined as Python functions, and then the **`verify`** function is called with these as parameters:

```
def C(): return read.PRESSURE >= 300
def R(): print("Pressure less than 300 PSI")

verify(C,R)
```

Note that due to Python's nameless **`lambda`** functions (also known as "lambda abstractions"), one can write the same call of **`verify`** as follows, without defining the condition as a separate function:

```
verify(lambda: read.PRESSURE >= 300,R)
```

7

[4] This is implemented by overriding the Python function `__getattr__(self,name)` that is part of any object and which gets called on an attribute (converted to a string) when the attribute is referred to but not found in the object .

verify_within(C,D,[R],[S])

This blocking construct verifies that the condition C **eventually** becomes true within the time duration D. The reaction R indicates the response that should be taken if the condition does not become true within the time duration. If the condition becomes true within its duration, then the function returns immediately without waiting for the duration to expire. The function returns true or false depending on whether the property succeeds verification or not. Assuming the condition and reaction from the previous example, the following example specifies a test that the condition becomes true within 10 seconds:

```
verify_within(C,10,R)
```

verify_within_voting(N,C_list,D_list,
[R],[S])

The argument C_list is a list of conditions and D_list is a list of durations. This construct verifies that at least N of the conditions in the condition list become true within their corresponding time durations in the duration list. The reaction R indicates the response that should be taken if this verification fails. If enough conditions evaluate to false (such that N conditions cannot possibly evaluate to true), the construct will execute the reaction immediately. As an example, consider three conditions C_1, C_2 and C_3, as well as a reaction R have been defined:

```
def C1(): return read.PRESSURE1 >= 300
def C2(): return read.PRESSURE2 >= 350
def C3(): return read.PRESSURE3 >= 375
def R():  print ("Pressure checks failed")
```

The following call verifies that at least two of the following conditions become true: C_1 within 5 seconds, C_2 within 10 seconds, or C_3 within 15 seconds; if this check fails, the reaction R is executed:

```
verify_within_voting(2,[C1,C2,C3],
                     [5,10,15],R)
```

assert_constraint(S,C,R,[D],[F])

This non-blocking construct registers the condition C to be continuously monitored "in the background", after which control returns immediately to the calling application. If the condition at some point evaluates to false, the specified reaction is executed in a separate thread, without disrupting execution of the calling application. If the flag F is true (or is not provided), then the monitoring continues after failure is detected, until expiration of D if provided. If a duration D is provided, the constraint is automatically removed after that duration. The constraint can also be removed explicitly with the function **remove_constraint(**S**)**. As an example, consider the following program text:

```
assert_constraint("C_One",C1,R1,25)
verify_within("C_Two",C2,10,R2)
```

This code will register the constraint C_1 to be monitored in the background, and will then immediately continue execution of the **verify_within** call. The condition C_2 will now be checked to become true *within* 10 seconds, while the condition C_1 is continuously being checked to be true *during* the 25 seconds. If constraint C_1 ever evaluates to false during the 25-second interval, reaction R_2 gets executed on a separate thread while the checking of constraint C_1 continues (because the optional reactivation flag F defaults to true), until the 25 seconds have elapsed and the monitor is removed.

conditional_interrupt(S,C,R,[D],[F])

This non-blocking construct is a variant of the former **assert_constraint** construct in that it registers the condition to be continuously monitored in the background, while control returns immediately to the calling application. However, if the condition at some point evaluates to true, the calling application is interrupted (temporarily stopped) while the reaction is executed. The monitor (also referred to as an interrupt handler) is by default removed if a provided duration D expires, or after the reaction is executed (unless the flag F is provided and is true). The interrupt handler can also be explicitly removed with the function **remove_interrupt**(S). The library also contains a function **timed_interrupt**(S,D,R), which interrupts the calling application after the specified duration and executes the specified reaction.

Controlling Functions

The LCS DSL library also contains functions for submitting information into the system, such as explicitly updating derived ("application-controlled") measurements, sending commands to end items, sending messages to displays, and initiating user dialogs via the displays. Two functions exist for issuing derived measurements from an application, enabling different applications to communicate via shared variables. The function **set**(S,E) submits a single update of the measurement named by the string S to the value returned by the expression E (where E represents an expression involving one or more measurements). The function **derive(**S,E**)** defines a new measurement named S to be derived from other measurements: querying S will henceforth return the value of expression E. This means that when any measurement referred to in E is updated, S gets updated to denote the new value of E. As an

example, the following code snippet will ideally print the value 500:

```
derive(X,lambda: Y * 100)
set(Y,5)
print X
```

Note, however, that the **set** and **derive** functions are not guaranteed to have immediate effect since they invoke the publish/subscribe middleware services. These constructs send a message to update the measurement, rather than instantaneously assigning a value to a local variable as in a traditional programming language. Hence, in the above example, some time delay would have to be inserted before the print statement to ensure that 500 gets printed.

The **send_command**(K) function submits a command object K, sub-classing a pre-defined Command class in the library. This is one of the few places where the user has to be aware of object-oriented features. A command's destination is defined and fixed in the LCS information model and cannot be changed by an application. Discussion of the LCS information model is beyond the scope of this paper, but it can be thought of as a database providing system configuration information. As an example, the following statement defines a reaction function that, when applied, submits a command to open a valve. The command is a so-called discrete command requesting an end item to assume a value from a discrete value space {OPEN,CLOSE}. The reaction function occurs as a parameter to a call of the **verify_within** function that checks whether the valve opens within 10 seconds, and commands the valve open in case it is not observed to happen.

```
open_cmd = Discrete("valve42",OPEN)

def R(): send_command(open_cmd)

verify_within(
  lambda: read.valve42 == OPEN,10,R
)
```

Finally, the DSL defines functions for sending messages to displays:

```
send_message(text,role,[criticality])
```

and initiating user dialogs via displays:

```
prompt(type,text,role,[listener_role],
       [timeout],[responses])
```

Such messages are targeted to subscribers (displays) that assume certain roles, hence following the publish/subscribe concept. A function also exists for spawning a new Python program, for either blocking or non-blocking execution:

```
perform(name,args,blocking)
```

Tabular DSL

One concern with using a DSL that includes the richness of a general-purpose programming language like Python is the potential for introducing complexity (and thus greater potential for errors) into control applications. This is of particular concern considering that LCS application developers are intended to be systems engineers. For this reason, LCS proof-of-concept team members at KSC developed a prototype Tabular DSL that was interpreted using the Python DSL as execution engine [6]. A tabular program consists of a sequence of lines, each line representing a DSL construct. The first column defines the construct, and subsequent columns define arguments. The tabular approach allows a customized editor to check input arguments and otherwise suggest alternatives, much like a syntax-directed editor.

5. PROTOTYPE IMPLEMENTATION

The prototype execution environment for the LCS DSL is implemented as a layered set of libraries in Python and C++, as illustrated in Figure 2. The messaging services were provided by a commercial implementation of the Data Distribution Service (DDS) [4]. The common services layer adapts the messaging services for the particular message data structures used in the LCS system. In keeping with the decision to adopt an internal DSL, the logic implementing the various language constructs described in Section 4 is coded in Python. This section presents the implementation of the common services and DSL layers of the LCS architecture, and discusses some interesting multi-threading issues that were encountered during the implementation and testing of these layers.

The C++ Common Services

As shown in Figure 2, the LCS architecture provides a set of common services for use by supervisory control applications and the command and telemetry gateways. The common services layer is implemented as a set of C++ API libraries.[5] Figure 3 presents a more detailed view of the common services design, including the interfaces and mechanisms it provides between the DSL layer and the DDS middleware.

A defining feature of the DDS messaging service is that it must be adapted for the specific message data structures it will manage. This is accomplished using tools provided by the vendor that compile message structure specifications defined in an Interface Definition Language (IDL) into C++ source code that will then implement the message class along with associated marshalling support and interface adapters. DDS defines a message topic in terms of one of these specific data structures.

[5] A complementary set of services were implemented in Java to interface with the display applications, as shown in Figure 2.

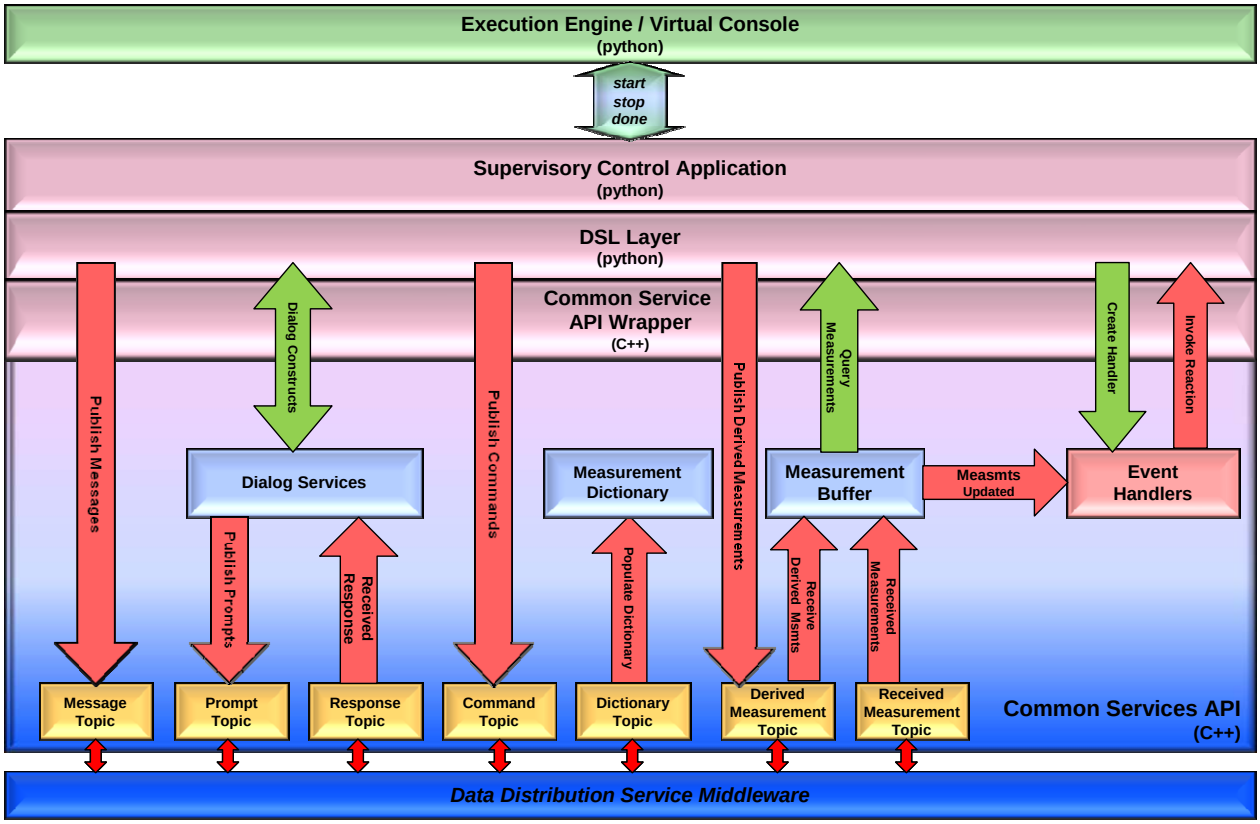

Figure 3. Common services design

To isolate some of the DDS-specific implementation details from higher-level libraries and applications, the C++ common services libraries also define their own distinct classes for each message type. This was found to be necessary because the DDS IDL translation process does not support the definition of methods in the message types – it interprets them as pure data structures. From a language design perspective, it was important to be able to support type-specific operations (e.g., comparisons) as member methods. So the C++ common services libraries provide interface adapters for each message type, which translate between the message classes used at the Python interface, and corresponding message structures generated from the IDL code used at the DDS interface.

In addition, a few other services are implemented in this layer for convenience. For example, since DDS messaging is asynchronous, the dialog service (see Figure 3) provides a single blocking dialog method that issues a prompt message and waits for a matching response. Implementing this service here helped to minimize the complexity of the interface between C++ and Python.

The measurement service provides a query interface by which the application may request the "current" (or most recently received) measurement sample for a named sensor (this interface is invoked through a **`read`**`.measurement` call in a DSL application). Because of the potential race condition between measurement samples arriving and when the application queries for samples, an additional measurement event service was required to enable the implementation of monitoring logic that would be sure to see every sample in a given sensor stream. The event service allows the application to register a callback handler that will be used to deliver every new sample associated with the named sensor as it becomes available. This mechanism is used in the implementation of more complex monitor functions in the DSL, including **`verify_within`**, **`verify_within_voting`**, **`assert_constraint`**, and **`conditional_interrupt`**. The measurement dictionary service provides a dictionary of known measurement (sensor) names and value types needed to allow the event service to verify requests before any data samples have been received.

The interface between C++ and Python is implemented using the Boost library [7], which supports exporting C++ classes and interfaces into Python. One of the key benefits of Python as the implementation language for the prototype LCS DSL was the existence of a well-defined extension API, and several distinct tools supporting interfacing between Python and other C++ tools. In particular, some of these tools support the mapping between Python's object model and that of C++. Of the several tools evaluated, the

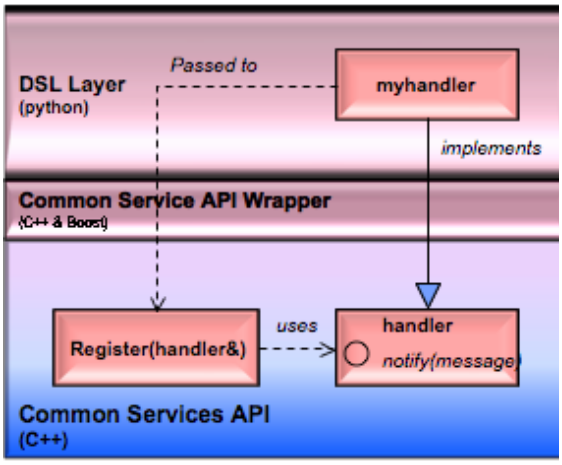

Figure 4. Event service design

Boost library (specifically, the Boost.Python library [8]) was found to provide the most complete support for symmetrical object-oriented interactions between C++ and Python. Other tools could export C++ functions or class methods into Python, but Boost was able to export a C++ interface into Python so that Python could implement a class implementing the interface that could then be called from the C++ layer. As will be shown, this interaction is key to the callback mechanism used to implement the event service.

Event services in the C++ common services layer are implemented using listener threads (usually one for each topic) whose job it is to notice when new messages have arrived in a given topic, match those messages against the set of sensor names for which notification callbacks have been requested, and then call each of the registered handlers, which are typically implemented in Python. This requires that these listener threads call through cross-language interfaces to deliver a data object that must also be translated from a C++ structure into a Python object.

To implement these interfaces, a C++ interface (abstract base class) defining the abstract callback method was exported to Python using Boost. This allowed the Python layer to then implement callback handlers derived from this interface, and to pass references to these to the callback registration methods which were also exported from C++ to Python (see Figure 4). Boost provides C++ templates that can be used to articulate an export wrapper for each class to be exported from C++ to Python. An export wrapper is a C++ class defined in terms of these Boost templates, which identify the Python-relevant properties of the class, and name the methods and members to be exported. The templates then implement all of the details needed to translate the C++ semantics and interfaces into the Python object model and API, which is implemented in C.

One key function that Boost did not provide, and that required significant additional work in the interface layer, was the synchronization of threads. The Python interpreter is designed to be thread safe, and the Python language can support thread safe operations (thread support can be dynamically enabled in a Python program so that single-threaded Python scripts don't pay the overhead cost of synchronization). Thread safety in the interpreter is achieved through the use of the Global Interpreter Lock (GIL), which is essentially a mutex lock on the entire execution engine. Since the threads used to implement the event callbacks originate in the C++ service library, and not inside of Python, these non-native threads need to interrupt the Python interpreter in order to call handlers in Python. Thus, additional code was required in the callback invocation methods to acquire the GIL before entering Python, and release it after returning. Apart from this detail, all of the interfacing details were handled by Boost templates.

The Python DSL

As already mentioned, the Python DSL is built on top of the C++ common services. A control application interacts with the common services though a set of objects of C++ classes, wrapped as Python objects/classes in the DSL layer. These include classes for creating objects containing various kinds of information, e.g., measurement objects. There are similar classes for creating command objects, message objects and dialog objects. In addition, the DSL layer includes classes that contain methods for publishing such objects to the middleware, that is, methods for publishing measurements, commands, messages, and dialogs. For example[6], the measurement service class offers three methods (amongst others):

```
class MeasurementService:
    def publish(name, value)
    def addHandler(name, handler)
    def removeHandler(name, handler)
    …
```

The method `publish(name,value)` publishes the value as a new measurement value for the measurement with that name. The two other methods are used for associating and un-associating measurement names with *handlers* that will get executed when the measurements are updated, typically due to state changes in the end items. A handler is an object of a class that subclasses the following `Handler` class:

```
class Handler:
    def update(measurement)
```

The `update` method of a handler object will execute code specific for a particular condition to be verified. Registering a handler with a specific measurement name causes the

[6] The code examples presented are abstractions of the real code. In some cases, Python syntax is not strictly followed, in order to make the presentation more succinct.

measurement service to record this association, and the handler will then get executed whenever that measurement is updated, for example when a new measurement arrives from an end item. Since the control application can execute while new measurements arrive in the measurement service, the control application and the measurement service execute in separate threads, resulting in a multi-threaded program. The measurement service thread reacts to the publishing of measurements on the message bus and activates handlers if required. In parallel, the control application thread reads measurements, publishes measurements, and registers and removes handlers.

The most interesting DSL functions are the monitoring functions that verify some property over a time interval, and which therefore need to register corresponding handlers with the measurement service. The handlers are all objects of a class `Constraint`, which subclasses the `Handler` class above, and hence overrides the update method. The `Constraint` class specifically contains the following data beyond the `update` method:

- *measurement name list* : list of measurement names referred to in the condition.
- *condition* : condition to be continuously monitored.
- *condition reaction* : reaction to be executed when the condition gets violated (becomes false).
- *duration* : duration over which the condition must be checked.
- *duration reaction* : reaction to be executed after this duration has expired.
- *reactivate flag* : true if the constraint should persist after a condition violation while the duration time period has not yet expired.

The `update` method has the following form:

```
def update(self, measurement):
  if not self.condition():
    spawn(self.condition_reaction)
```

The method is here somewhat simplified, in particular code has been left out which manipulates a state machine in order to ensure a correct interaction with the associated timer thread when it comes to removing the constraint. This (somewhat simplified) `update` function just checks the condition and spawns the reaction in a new thread in case it is violated.

In the measurement service class a method exists for adding a constraint:[7]

```
def addConstraint(C,R_C,D,R_D,F):
  constraint = Constraint(C,R_C,D,R_D,F)
  for m in measurements(C):
    ms.addHandler(m,constraint)
  constraint.startTimer()
```

The method creates a constraint from its arguments, and registers the constraint with each measurement in the measurement list (i.e., creates a handler for each measurement in the measurement list). This list is computed from the condition being evaluated and contains all measurement names mentioned in the condition. In the real implementation this measurement list is, however, not extracted automatically as shown. Instead this list is given explicitly as argument to all DSL functions taking a condition as argument. Automatically extracting the list of measurement names requires parsing of Python expressions. This will be discussed in more detail in Section 7 of the paper. Registration of the constraint has the following effect: *"check condition `C` during the period `D`; execute reaction R_C if `C` gets violated; execute reaction R_D if duration `D` expires; repeat verification on condition failure during the period if `F` is true"*.

The LCS DSL's monitoring functions are all defined using a generic non-blocking function **`monitor`**, which is based on the functionality of `addConstraint`. A call of the form:

```
monitor(C,R_C,D,R_D,F)
```

invokes `addConstraint(C,`R_C`,D,`R_D`,F)`, after first having evaluated the condition to check its immediate value. Note that this function is given a reaction to execute when the condition is violated as well as a reaction to execute when the timer expires. After registration, the function terminates and program execution continues, with the constraint now active in the background.

The specific monitoring functions can now be implemented as follows. The translation of the **`assert_constraint`** function is straightforward:

```
def assert_constraint(C,R_C,D,F):
  monitor(C,R_C,D,None,F)
```

A constraint will be registered, which checks the condition `C` continuously in the background (relative to the control application thread), and executes the reaction R_C if it gets violated. The constraint condition is monitored until the duration `D` expires. No special reaction is executed when the time period expires beyond removing the constraint (this default timeout behavior is built into the **`monitor`** function).

The **`conditional_interrupt`** function behaves similarly to **`assert_constraint`**, but differs in two significant ways:

1. <u>*condition negation*</u> : the condition passed as argument triggers the reaction to be executed when it becomes *true* (in contrast to false).

[7] Note that the Python-specific `self' argument is omitted here for the sake of clarity in the presentation.

2. *interrupt* : The reaction to be executed has to interrupt and stop the main DSL control application while executing.

The first objective is simply achieved by negating the condition that is stored in the constraint. The second objective is achieved by letting the control application thread execute these interrupt reactions, which are stored in a global queue with the function **addInterrupt**:

```
def conditional_interrupt(C,R_C,D,F):
  R_I = lambda: addInterrupt(R_C)
  monitor(negate(C),R_I,D,None,F)
```

As can be observed, the reaction to be executed when the negated condition goes false (i.e., when the original condition goes true) is a reaction that adds the argument reaction to the interrupt queue. No special reaction is executed when the time period expires beyond removing the constraint. Interrupts are now executed by a special function **executeInterrupts()**, which is called as the first statement in all DSL functions. This function iterates through the interrupt queue, and executes them one by one. Note that in this manner only DSL functions (and not the basic Python constructs) can be interrupted, so there may be small delays in seeing the main thread get interrupted, until the main thread begins execution of the next DSL construct in the flow of the application.

Timed interrupts share the interrupt characteristic with conditional interrupts (they stop the main thread), but they are purely triggered by the expiration of a timer. Technically, a timer (Python's library provides timers as threads) is started that triggers after the duration expires and executes in a separate thread a reaction that inserts the interrupt in the interrupt queue. This is achieved with the following function, requesting the reaction R_D to get executed as an interrupt after D time units:

```
def timed_interrupt(D,R_D):
  R_I = lambda: addInterrupt(R_D)
  monitor(None,None,D,R_I,False)
```

The call of **monitor** has a `None' as condition (meaning: no condition), a `None' as condition-reaction R_C, the duration D, and a reaction R_I that adds the argument duration-reaction R_D to the interrupt queue. The reactivate flag is false, which in fact is irrelevant since no condition is monitored.

The semantics of the **verify_within**(C,D,R_D) construct is to block until the condition C becomes true, or until the duration D expires, at which point R_D gets executed. In order to support the blocking wait on the condition evaluation (which occurs in a separate thread), a new class ConstraintStatus is introduced. Objects of this class serve to establish communication between the caller and the condition evaluator. The class has the following interface:

```
class ConstraintStatus:
  def wait()
  def signalSuccess()
  def signalFailure()
  def satisfied()
```

An object of this class contains a Boolean flag indicating whether the condition has been satisfied or not. The function signalSuccess sets this flag to true and signals a semaphore that the wait function is "on". The signalFailure function sets the flag to false and signals the semaphore. The function satisfied thereafter returns true in the first case and false in the second case (the value of the satisfaction flag). To implement the **verify_within** construct, the functions signalSuccess and signalFailure are passed as arguments to the generic **monitor** function, representing condition reaction and duration reaction respectively, as follows (note also that the condition is negated to model the "eventually true" check):

```
def verify_within(C,D,R_D):
  cs = ConstraintStatus()
  monitor(
    negate(C),cs.signalSuccess,
 D,cs.signalFailure,False
  )
  cs.wait()
  if cs.satisfied():
    return TrueCode
  else:
    R_D
    return FalseCode
```

The implementation of the **verify_within_voting** function is very similar to that of the **verify_within** function. However, it uses a slightly more complicated variant of the ConstraintStatus class with the same public interface. The new class contains two additional counters: the number of conditions that minimally must be satisfied, and the number of candidate conditions that have not yet evaluated true. The class constructor is called with two natural numbers: the number of conditions that minimally must evaluate to true, and the number of initial candidate conditions. The function signalSuccess is called when a condition becomes true within its time period and it decrements the number of conditions required to still become true, and also the number of candidate conditions. Similarly, the function signalFailure is called when a time period for a condition expires without the condition having become true, in which case only the number of candidate conditions is decremented. The function satisfied() returns true if the number of required conditions is non-positive. For each condition-duration pair a constraint is registered which calls signalSuccess when the condition becomes true (it monitors the negation being true and reacts when the negation becomes false), and which calls signalFailure when the time period expires without the condition having become true.

Multi-Threading Issues

As outlined above, the implementation is multi-threaded due to the fact that constraints, reactions and timers execute in parallel with the main application. This raises a number of issues. One issue is the general notion of program suspension, as required by interrupts, external suspension commands (the **`pause`**() construct, not discussed in Section 4) and hard program termination (normal as well as abnormal). Our approach to interrupts is to let them execute on the main thread as described above: the main thread executes queued interrupts. In a single-threaded application, program termination simply means terminating that thread in a safe place, after cleaning up various data structures. If the termination is abnormal and deep inside a nested series of function calls, throwing an exception (after cleanup) may be the most convenient manner in which to terminate the program. In the presence of multi-threading, however, the situation becomes more complicated. When a thread must terminate the entire application, it needs to inform the other threads about this intention. Note that throwing an exception in one thread does not terminate the other threads: exceptions are local to threads. Furthermore, a terminated thread must not just clean up and delete shared data structures, such as the common services. When the main application terminates, what normally happens is that all common services are closed down. However, if there are still threads running (e.g., other applications), one must wait for these to terminate. Design of a more robust and satisfactory long-term solution to this problem was left as future work for the LCS team.

A second issue is data races and deadlocks. As part of a unit testing framework, an algorithm was implemented for detecting cycles in lock graphs from normal (non-deadlocking) runs [9]. This became highly useful for ensuring the lack of deadlock potential. One particularly interesting deadlock occurred due to the interaction between the Python interpreter and C++. The Python interpreter takes what is called the Global Interpreter Lock (GIL) on the entire virtual machine each time it executes a statement. The unexpected deadlock scenario was as follows: the Python program executes a statement (takes the GIL) that attempts to take a lock *L* on the C++ side; in parallel, a thread on the C++ side takes the same lock *L* and then makes a call-back of a Python function (a condition evaluator), causing the GIL to be requested. Now the Python thread holds GIL and the C++ thread holds *L* – a deadlock has occurred. To resolve this potential deadlock problem, additional code was required in the measurement service callback invocation methods to acquire this lock before entering Python, and release it after returning.

A third issue is the removal of constraints. A timer can expire while a constraint reaction is being executed, in which case the timer and the reaction have to agree who removes the constraint. In the current design, a shared state machine keeps track of occurred events and hands removal permission to the right thread.

A fourth issue is a dangling pointer problem that arises when a Python thread passes a Python object reference from the Python data area over to the C++ data area, and then later exits the scope where that object is introduced. In this case the object gets garbage-collected on the Python side, and the C++ side now holds a dangling pointer (pointing to a garbage-collected data area). This issue was resolved by always making sure on the Python side that such objects were not garbage-collected, by inserting them into a list of "garbage objects" (although certainly not garbage on the C++ side).

6. EVALUATION

An evaluation was performed of the implemented monitor and control system. The monitoring and control system consists of the Python DSL implemented in Python, the underlying common services layer implemented in C++, an information architecture (database) containing persistent monitor and control meta-data, display and GUI software implemented in Java, and the two COTS tools providing the gateway and publish/subscribe middleware functions. In addition, the evaluation also involved writing control applications in the Tabular DSL.

The Case Study

The monitoring and control system was evaluated by application to two different scenarios for the existing Space Shuttle system, both part of the current Shuttle launch countdown sequence: (i) a Liquid Hydrogen Fast Fill application ("LH2"), and (ii) a Main Propulsion System pneumatic decay test ("MPS"). A high-fidelity software simulator replaced the real Shuttle end item during this evaluation.

The LH2 scenario consists of monitoring the filling of the Space Shuttle external fuel tank with liquid hydrogen, which is transported from a container positioned in some distance from the Shuttle and connected via a complicated system of pipes, valves, etc. The LH2 process begins at launch minus 6 hours and 50 minutes and lasts for about 40 minutes, nominally. Its purpose is to fill the shuttle's LH2 external tank from a level of 5% full to 98% full. During that time liquid oxygen is also transported to the Shuttle (called the LO2 process). Sample code representative of a small portion of the LH2 control application is shown in Figure 5. The MPS process measures the pressure decay and flow rates in the main propulsion system's LH2 and LO2 pneumatic system lines. The MPS application was implemented to demonstrate to the Shuttle operations community that the DSL could be used to implement applications with more algorithmic and numeric computations than the LH2 application.

```
# A3301 PRIMARY OPEN COMMAND
send_command( discrete( "VALVE1", "ON" ) )

# TEST INDICATORS
if verify_within_voting ( 2, [lambda : read.INDICATOR1 == OFF, # VLV1 CLOSED #1 IND
                              lambda : read.INDICATOR2 == ON], # VLV1 OPEN #1 IND
                              ["INDICATOR1", "INDICATOR2"],
                              [8, 26], # 8 seconds, 26 seconds
                             DIALOG,
                             "INDICATOR1 and INDICATOR2"
                          ) > 0 :

                          # A FAILURE OCCURRED.  ERROR MESSAGE
                          send_message ( "VALVE1 PRIMARY COMMAND FAILED", LH2, WARNING )

                          # USE SECONDARY COMMAND AND SECONDARY INDICATORS
                          perform( "VLV1_SEC_OPEN", [], BLOCKING)

# SUCCESS MESSAGE
send_message ( "VALVE1 OPENED SUCCESSFULLY", LH2, INFORMATION )
```

Figure 5. Sample DSL code to open a valve and check for appropriate measurements within specified times.

During the tests, a monitoring display showed the current state of affairs on the (simulated) launch site. Information was provided to this display application (written in Java) via subscriptions on the message bus, and directives from the operators (e.g., initiation of the LH2 application) were input into the LCS system via publication on the message bus. A snapshot of the LH2 display is shown in Figure 6, showing the liquid hydrogen storage tank on the left, connected to the Space Shuttle external tank on the right.

The DSL applications send sequenced commands, verify their successful execution by checking specific measurements, send messages and prompts to operators on displays like the one in Figure 6, monitor measurements for condition violations, and automatically take actions in response to such violations.

The DSL applications were developed using the Eclipse PyDev development environment [10]. The programs were designed to meet requirements and specifications provided by Shuttle system engineers. The control applications were also reviewed by the system engineers to assess their correctness, completeness, and readability.

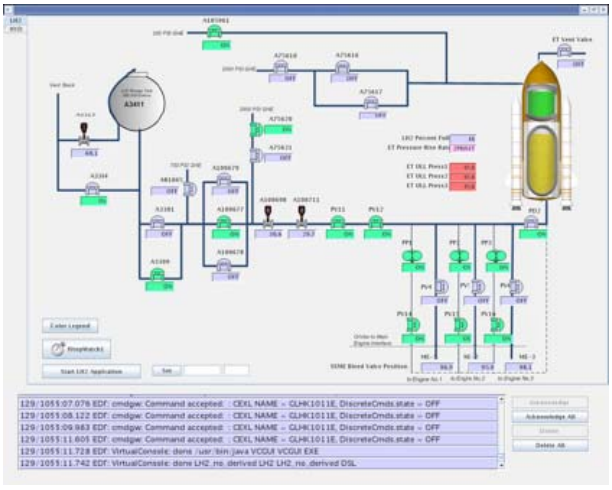

Figure 6. Display for the LH2 application tests.

Result of Evaluation

Implementation of the control applications in the Python DSL was relatively straightforward. Eclipse and the PyDev development environment were useful in quickly turning around code changes and bug fixes. Application software developed using the Python DSL successfully executed the LH2 scenario, demonstrating the expected functionality and satisfying the requirements for the proof-of-concept activity.

One significant V&V challenge was that there was no way to actually run and test the application programs without having the full LCS system and simulator up and running, or without intermingling supervisory control code with simple simulation “stubs” within the control application. This highlights the general need for compositional software development and unit testing, where each component can be tested in isolation. It furthermore illustrates the complication resulting from integration of several software artifacts either written in different languages or written by different parties.

Since Python is dynamically typed, the static checkers PyLint [11] and PyChecker [12] were used to statically check Python programs before their execution. None of these checkers, however, perform type checking, and only catch some of the errors that would be otherwise caught in a statically typed programming language. To catch type errors, one could potentially write a specialized static checker for the Python DSL. Python’s meta-programming features enable the development of specialized checker tools for Python. The alternative is to use testing techniques, in which case a coverage tool like Coverage [13] seems to be crucial in constructing good test suites, in combination

with a unit testing tool like PyUnit [14]. Based on the LCS proof-of-concept team's experience, these are all useful tools, but testing the program by running it still provides the best chance of finding type errors in Python. For this purpose, simulators and stubs for other software components in the architecture are of critical importance.

In summary, Python proved to be a convenient and functionally effective language for the purposes of implementing the monitor and control DSL for the presented scenarios. However, its dynamic typing presents a serious challenge from a verification point of view. Further analysis would be required to determine whether development risk can be adequately reduced, perhaps through the use of a specialized IDE that guides application developers and prevents them from making certain classes of errors.

7. DISCUSSION

Python is clearly an easy-to-learn and convenient language, due to its succinct syntax, abstract constructs and extensive library, including meta-programming features. It supports object-oriented as well as a limited form of functional programming. In particular, the special lambda-expression allows one to write anonymous (nameless) functions, which are convenient when implementing a DSL. Python provided a powerful tool for the short-fuse and resource-constrained LCS proof-of-concept activity. In addition, the Boost.Python interface library proved to be extremely useful for implementing the interfaces between Python and the C++ messaging frameworks.

The LH2 example application demonstrated that the required control idioms could be implemented in a programming language interface to the new distributed LCS architecture. The program was able to express both the sequential logic needed to implement test procedures (issuing commands, and verifying responses through sensor measurements), and the reactive logic needed to implement safety constraints. It further demonstrated the use of multiple programming languages specialized to support the implementation of the specific control idioms, and user interfaces.

The first issue of concern is the lack of static typing in the language. Python is dynamically typed, meaning that many kinds of errors are not caught at compilation time. There do exist static checker tools for Python, such as PyChecker and Pylint, but they do not seem to catch quite fundamental problems, such as, for example, the wrong number of arguments in a function call. The Tabular DSL demonstrated some mitigation of these risks. Although the risks were not entirely eliminated in the proof-of-concept task, it did demonstrate that tools were available to support the implementation of this verification at several points in the process, including in the language, or in the integrated development environment, or both. These choices are left for the implementing project.

The second issue of concern identified during the initial DSL survey was the uncertainty associated with the ability of Python to satisfy performance requirements. The ability of the applications to meet real-time deadlines – primarily, the latency of reactions to changes in sensor values – will depend to a large extent on the latency of the delivery of the measurements to a gateway, and from there across the message bus to the monitoring application. That is to say, the performance issue is more of an architectural concern for LCS, considering the adoption of a distributed architecture. With this said, the language allows applications to react to measurement changes as they arrive via the callback interface; there will be a maximum rate at which these update events can be delivered and processed that will depend on how many conditional expressions the application uses. It is not expected that the post-proof-of-concept DSL and common services implementation would represent a performance bottleneck in the LCS system, but formal benchmarking remains to be performed.

One of the decision points in the design of the DSL API was whether it should have an object-oriented flavor in the spirit of Python, or whether it should have a more classical procedural flavor like the C programming language. Python allows both forms. In the object-oriented approach, the DSL programmer would have to be aware of classes and objects and would work with dot-notation. In the procedural approach, a program consists of a sequence of function definitions. It was decided that the object-orientation of Python should be downplayed in order to minimize the number of Python concepts a system engineer would have to master: stand-alone functions resemble the syntactic constructs in familiar programming languages, and may also lead to more succinct control application implementations. This decision was largely influenced by the desire for the new language to reflect the capabilities of the GOAL language used as a requirements model, and to make the language intuitive to the system engineers who would be using it. The resulting syntax, however, exhibits idioms of both programming styles, and this can be seen as a potential source of programmer confusion.

Several of the constructs take as parameters a lambda-expression (or a function reference) as condition and are supposed to evaluate this condition each time a variable occurring in that expression (or referenced by that function) is updated. This concerns the constructs **`verify_within`**, **`verify_within_voting`**, **`assert_constraint`**, and **`conditional_interrupt`**. These constructs cause update handlers to be registered with common services and take measurement name lists as explicit arguments so that update handlers can be registered for those measurement names. These measurement name list arguments are an annoyance and would preferably be inferred automatically

from the conditions themselves. This would require parsing and extraction of variable names from the condition expressions (lambda-expressions or functions). This could be non-trivial if the conditions contain nested function calls, leading to arbitrarily "deep" Python code, in which case a free variable analysis of all of Python is required. There are currently no obvious simple solutions to this problem. Alternatively, one could specify a way of writing conditions where a condition is not a general Python expression. For now, the design assumes that the user specifies the list of measurements, any update of which should trigger the check of the condition. In the future, condition expressions could be parsed at least one level deep.

There is currently an asymmetry in the design of the DSL from a linguistic point of view. For example, consider the blocking function **`verify_within`** and the non-blocking **`assert_constraint`**: **`verify_within`** checks that a condition becomes true *within* a time period, whereas **`assert_constraint`** checks that a condition stays true *during* a time period. To be completely symmetric, the language would need to offer additional constructs, such as a blocking **`verify_during`** function, which checks that a condition stays true during a time period, and a non-blocking **`assert_constraint_within`** that checks that a condition becomes true within a time period. The prototype DSL design was driven by requirements provided by experts at KSC in a very direct manner, hence the asymmetric design. Such a requirements-driven design might subsequently be followed by a "linguistic cleansing" procedure, introducing symmetry as appropriate, and providing other desirable language characteristics.

In general, the DSL can be regarded as a temporal logic embedded into an imperative programming language, a concept that can be given further theoretical attention. The **`monitor`** function described in the implementation section in fact represents a very general temporal operator that could be embedded into any programming language.

A final issue to be discussed is the notion of time. In the current DSL implementation, time is measured on the clock local to the control application. An alternative would be to measure time in terms of time stamps in measurements, or some other global time-keeping mechanism. This issue is of particular importance given the distributed nature of the LCS architecture.

8. CONCLUSIONS

Experience with the current Space Shuttle test and checkout system suggests that significant lifecycle costs are associated with the development and maintenance of software for monitor and control applications. Particularly expensive is the process by which system engineers express requirements for test procedures in prose, software developers translate these requirements into code, and then both sets of experts are engaged in verification of the resulting application's correctness. The Constellation Launch Control System will have to significantly reduce these lifecycle costs while assuring a consistent high level of safety and security.

To this end, the LCS project has completed a proof-of-concept demonstration of an approach that has the potential to ultimately lower the lifecycle costs of monitor and control applications for launch processing, test and checkout. This paper has presented one of the important contributions from this proof-of-concept activity, namely the "home-grown" Python DSL and an associated common services framework design.

In summary, the use of Python as the base language for the "home-grown" LCS DSL was substantially successful. The prototype demonstrated that the implemented DSL design could satisfy the requirements of a realistic monitor and control problem, working effectively with the other elements of the distributed LCS architecture. The KSC system engineers readily understood the DSL logic, supporting its development throughout the prototype. Along with Python's strengths as a popular, well-supported language, the Python-based DSL seems to live up to its promise as a low-cost solution. However, given the paramount importance of safety and reliability, Python's weak static typing remains a concern.

ACKNOWLEDGEMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The authors would like to acknowledge the contributions of the rest of the LCS Proof-of-Concept team at NASA Kennedy Space Center, and thank James Shaver for his feedback on a draft of this paper.

REFERENCES

[1] The White House, Office of the President, "A Renewed Spirit of Discovery: The President's Vision for U.S. Space Exploration," Released to accompany the President's NASA FY 2005 Budget, January 2004; URL: http://www.nasa.gov/pdf/55583main_vision_space_exploration2.pdf

[2] Constellation Program web site; URL: http://www.nasa.gov/mission_pages/exploration/main/index.html

[3] T. Mitchell, “A Standard Test Language – GOAL (Ground Operations Aerospace Language)”, Proceedings of the 10th Workshop on Design Automation at the Annual ACM IEEE Design Automation Conference, pp. 87-96, 1973.

[4] “Data Distribution Service for Real-time Systems, v1.2”, Standard published the Object Management Group, Inc., January, 2007; URL: http://www.omg.org/technology/documents/formal/data_distribution.htm

[5] Python web site; URL: http://python.org

[6] K.W. Leucht, G.S. Semmel, "Automated Translation of Safety Critical Application Software Specifications into PLC Ladder Logic", Proceedings of the 2008 IEEE Aerospace Conference, Big Sky, MT, May 2008.

[7] Boost web site; URL: http://boost.org

[8] Boost.Python web site; URL: http://boost.org/libs/python/doc/index.html

[9] J. Harrow, “Runtime Checking of Multithreaded Applications with Visual Threads, SPIN Model Checking and Software Verification”, Lecture Notes in Computer Science Vol. 1885, Springer, pp. 331-342, 2000.

[10] PyDev web site; URL: http://pydev.sourceforge.net

[11] Pylint web site; URL: http://www.logilab.org/857

[12] PyChecker web site; URL: http://pychecker.sourceforge.net

[13] Coverage web site; URL: http://www.nedbatchelder.com/code/modules/coverage.html

[14] PyUnit web site; URL: http://pyunit.sourceforge.net

BIOGRAPHIES

Matthew Bennett *is a Senior Software Systems Engineer in the Flight Software & Data Systems section at the Jet Propulsion Laboratory. He has interests in model-based engineering, software architecture, and spacecraft autonomy. He has developed mission software for fault protection, guidance and control, science data collection, performance analysis, and simulation. He holds an MS from the University of Washington in Computer Science, and a BS from the University of California at San Diego in Computer Engineering.*

Richard Borgen *is a senior software engineer in the flight software applications and data management group at the Jet Propulsion Laboratory. His specialties include database management and ground telemetry systems. He received his BSEE from Michigan State University, and his MS in Systems Management from the University of Southern California.*

Dr. Klaus Havelund *is a principal computer scientist in the flight software applications and data management group at the Jet Propulsion Laboratory. His research interests include verification of software wrt. models, in particular using dynamic analysis techniques such as runtime monitoring. He received his BS, MS and Ph.D degrees in Computer Science from the University of Copenhagen, Denmark. The Ph.D was executed at Ecole Normale Supérieure, Paris, France.*

Dr. Michel Ingham *is a member of the technical staff in the Flight Software Systems Engineering and Architecture Group at the Jet Propulsion Laboratory. His research interests include model-based methods for systems and software engineering, software architectures, and spacecraft autonomy. He earned his Sc.D. and S.M. degrees from MIT in Aeronautics and Astronautics, and a B.Eng. in Honours Mechanical Engineering from McGill University in Montréal, Canada.*

David Wagner *is a senior software system engineer in the flight software applications group at the Jet Propulsion Laboratory. He has a BS in Aerospace Engineering from the University of Cincinnati, and a MS in Aerospace Engineering from the University of Southern California.*