

JavaTM magazine

By and for the Java community 

JAX-RS 39 | JAVA INTERFACES 47 | FANTOM 53

SEPTEMBER/OCTOBER 2016

Devices and IoT

14 INTERACTING
WITH SENSORS

24 CONTROLLING
CNC ROUTING
FROM A
RASPBERRY PI

32 IOT AND
THE CLOUD

Capturing and responding to data are the heart and soul of IoT. Here's how to do both with Java on an inexpensive board that uses an x86 processor.

01

XRebel

 ZEROTURNAROUND

THE LIGHTWEIGHT JAVA
PROFILER

TRY IT FREE NOW!

*Get a free
t-shirt! →*



//from the editor /



Appreciating Limited Choice in Languages

The more prescriptive a language is in the details, the easier it is to code productively.

Our coverage of JDK Enhancement Proposals (JEPs) in this issue examines a recent proposal to standardize the syntax of command-line arguments for tools that ship in the JDK. As the proposal points out in support of its core concern, presently there are multiple ways of asking for help from the command line. And if you happen to guess wrong when using a given tool, you need to circle through the variety of possibilities. These can vary from `-help` to `--help` to `-?`. And then there's the unmentioned last resort, which is to run the program with no arguments and see what kind of information you get in the error message.

I wholly support this standardization, but I'd go much further. In my view, the syntax of

command-line switches should be included in style guides for the language. If the Java team had specified a standard convention for switches when it released the language (in the same way that it recommended initial capitals for class names and all capitals for constants), this small annoyance would not exist. The more a language can formalize small details, the easier it is to get things done.

But in an ideal world, even this solution is insufficient. I strongly believe that the abundance of Java style guides is itself a limitation. I'd far prefer that there be one consistent set of recommendations that was universally followed. For example, writing out definitive guidelines for the

PHOTOGRAPH BY BOB ADLER/GETTY IMAGES

ORACLE®



Java in the Cloud

Oracle Cloud delivers high-performance and battle-tested platform and infrastructure services for the most demanding Java apps.

Oracle Cloud.
Built for modern app dev.
Built for you.

Start here:
developer.oracle.com

#developersrule



location of opening braces, size of indents, tabs vs. spaces, how to stagger or not stagger if/else sequences, Javadoc's numerous formatting options, and so forth. Obviously, this would apply to higher-level concerns as well: fully expanded imports vs. wildcards, the sequence of import statements and variable declarations, and so on. By having a fixed set of guidelines, every Java listing would be consistent and not require re-examination to adjust to a given individual's or site's style.

Curiously enough, in a certain way, Java's initial appearance on the scene addressed what, at the time, was a tremendous laxity in language that made some tasks exceedingly tedious. The principal language before Java was C. It was purposely designed from a radically nonprescriptive perspective. Even today, after many rounds of standardization, C has numerous places where behavior is undefined or left up to the implementation to define. In the mid-'90s when Java first appeared, C was far looser. An integer could be more or less anything the compiler defined it to be, with—if I recall correctly—a minimum of 16 bits of width. 16, 32, and 64 bits were all legitimate implementa-

tions of an integer. As a result, porting C from one platform to another was extraordinarily tedious. Java solved these problems. Data items had fixed sizes across platforms, and code could be run on multiple platforms without modification.

C's lack of standardization caused so much pointless activ-

Most modern development organizations prescribe their own “house style” for code, which is frequently enforced in code reviews. But these styles conflict with each other for lack of a single, unified set of conventions.

ity that when the original team from AT&T Bell Labs developed a new language, Go, they chose a highly prescriptive implementation. There is one formalized coding style for Go, and all code is expected to use that style. A code formatter is bundled with the Go distribution. In the language itself,

there are additional constraints. For example, the executable code after any if statement must be enclosed in braces, even if it contains only a single line. Many other conventional items are defined by what is known as “idiomatic Go.” The happy result is that all Go code looks the same. Reading and writing it is easy.

The lack of standardization of details during the last few years has been an issue in Java in small but annoying ways, beyond command-line syntax. For example, the three variants of the annotation for indicating a field should not be null: @NonNull, @Nonnull, and @NotNull. The first of these was used by Checkstyle and FindBugs, the last of them by Java EE 6 and the IntelliJ IDE. The result was that if you coded in one environment and moved to a different development environment, you had to change your code or your toolchain to get your expected level of null-checking. This is, of course, the exact antithesis of Java's vaunted portability, when switching IDEs is enough to make code behave differently. Fortunately, Java 8's use of the Checker framework has now consolidated the convention around @NotNull.

The convenience and benefits of such strictures that ensure uniform syntax are widely recognized. This is evident in the choice of most modern development organizations to prescribe their own “house style” for code, which is frequently enforced in code reviews, so that all developers use the same conventions. But these styles conflict with each other for lack of a single, unified set of conventions.

If the world were all Java, I think it would not be too onerous to put up with the differences, although the time lost in doing so is lost for no good reason.

But in an increasingly polyglot world in which other languages (JavaScript, HTML, and so on) play significant roles, the lack of enforced coding standards in Java and especially across those languages combine to create a sustained and pointless drag on productivity.

Andrew Binstock, Editor in Chief
javamag_us@oracle.com
[@platypusguy](https://twitter.com/platypusguy)

*Coming in November:
our special issue on JUnit 5.*



JULY/AUGUST 2016

The Limitations of JSON-P

The article by David Delabassée titled “Using the Java APIs for JSON Processing” in the July/August 2016 issue does not mention significant problems with the JSON-P standard, which is poorly designed, and the reference implementation, which is poorly implemented.

I believe the standard is poorly designed because neither of the two APIs in the standard support parsing and serialization of plain old Java objects (POJOs) to and from JSON, the natural fit for an object-oriented language.

JSON parsers and serializers written with the JSON-P API are too big, require too many API calls, and are too fragile, requiring extensive changes if the structure of the JSON to be processed changes.

I believe the reference implementation is poorly implemented because in my use case, converting thousands of Java objects to JSON strings of 500 to 50,000 characters each, a program using the Jackson API's object model converted a real data set to JSON an astounding 10 times faster than one using the object model API in the JSON-P reference implementation. The presumably more efficient streaming API in JSON-P would have required unmaintainable code in which the API calls to open JSON arrays and objects were often widely separated from the API calls to close them.

So, I believe you should discuss alternative JSON processing libraries such as Jackson and GSON.

—Jeffrey S. Mayo
Norcross, Georgia

David Delabassée responds: “In the article, I discuss marshaling: ‘Note that binding (that is, marshaling of Java objects to JSON documents and vice versa) will be addressed in a related API, the Java API for JSON Binding (JSON-B), which is currently being defined in JSR 367.’”

The editors add: “JSON-P is part of the Java EE standard and has an elegant parsing interface; in addition, it’s fast. When JSON-B is finalized, we’ll cover it in Java Magazine as well. Jackson and GSON both represent good choices for developers needing JSON serialization now.”

The Practice of Small Classes

With respect to your editorial in the July/August issue, “The Problem of Writing Small Classes,” I second your idea and argumentation about writing small classes in any “normal” code project. But when it comes to web development (servlets) or when writing a RESTful API, I always find it quite painful to externalize lots of pure (RESTful) logic just to stay within my own limits.

Just as an example, I have a REST resource class called `ItemResource`. I define several methods that make use of different media types and define several request mappings and request methods (using Spring, but it would be the same with pure Java EE).

So, alone with the imports, injections, method definitions, annotations, and logging plus the pure calls to business (Entity) classes and some Java 8 stream processing (mostly one-liners), I easily reach 200 lines of code or more.

I start to believe in the best of both worlds—trying to keep it simple and small but still binding semantic units together. A class should still be a class, and if we divide this more and more into smaller parts we could end up with one-method classes or endless delegations to deeper classes.

I tend to not limit myself to a hard number of lines, but take other measures in metrics plugins such as method complexity and try to take the separation of concerns seriously.

—Alex Hepp
Germany

We welcome comments, suggestions, grumbles, kudos, article proposals, and chocolate chip cookies. All but the last two might be edited for publication. Write to us at javamag_us@oracle.com. For other ways to reach us, see the last page of this issue.

ORACLE®

Coming this fall to selected Linux/**ARM** devices

AOT-Compiled Java for ARM

→ **Starring: Your Java app and the Excelsior JET Runtime**

From version 11.3, Excelsior JET will begin supporting Java SE 8 on ARM-based platforms

Get Your Early Access Copy Now

Registration-Free Download



PHOTOGRAPH BY BENBENW/FLICKR

10

```
//events /
```

Way to Build,” led by IBM’s Dan Heidinga.

Topconf

NOVEMBER 15, WORKSHOPS

NOVEMBER 16–18, CONFERENCE
TALLIN, ESTONIA

Java and the JVM, reactive architectures, sustainable development, and big data highlight this year's conference. Couchbase Developer Advocate Laurent Doguin will present a practical example of RxJava, and MATHEMA Software Senior Consultant Thomas Künneth will discuss current approaches to Java on mobile devices.

Java Enterprise Summit

NOVEMBER 28–30

BERLIN, GERMANY

What does a state-of-the-art enterprise Java application look like? Which APIs are useful? What are the roles of various web and JavaScript frameworks? And how important is standard Java EE today? These and many other questions will be discussed at this year's Java Enterprise Summit. A large training event is held concurrently with the Micro Services Summit, hosting 24 power workshops with well-known German

microservices and enterprise Java experts. (No English page available.)

ConFoo

DECEMBER 5-7

VANCOUVER, CANADA

This multitechnology conference for web developers features sessions on Java and JVM. Scheduled topics include Java 9; caching; machine-learning models with Java- and Spark-based tools; Docker and Java; and writing better streams with Java 8.

Special Note: Event Cancellation

QCon Rio

OCTOBER 5-7

RIO DE JANEIRO, BRAZIL

The organizers report: “Faced with an unstable political and economic environment . . . we considered it prudent to cancel our edition of QCon 2016. We emphasize that this decision does not affect the preparation of QCon São Paulo 2017.”

Have an upcoming conference you'd like to add to our listing? Send us a link and a description of your event four months in advance at javamag_us@oracle.com.

//event recap /

JVM Language Summit 2016

For the past nine years, Oracle (and earlier, Sun) has been sponsoring a small gathering of Java experts whose primary work is on the JVM. The roughly 150 attendees gather to discuss the JVM and JVM languages.

Unlike most conferences, at this summit there is only a single track, attended by all the participants. Over three days, many sessions of surpassing interest are presented. All of them are recorded. The videos for the 2016 summit are posted on [YouTube](#). This year's summit in early August saw particular focus on Java 9 and post-Java 9 releases.

As for JVM languages, you can find videos of sessions on implementation details and upcoming features for Kotlin and Scala. And there are sessions that focus on standalone technical topics: immutable collections, bytecodes, use of Intel vector instructions, and so forth.

While the sessions are technically demanding, you do not need deep knowledge of the JVM to follow along. However, as you'll see in the videos, you do need some fluency with the terminology of JVM functions and features.

The conference is held in July or August at Oracle's campus in Santa Clara, California. While attendance is intentionally kept low, developers working on JVM topics can apply to attend on the conference's [home page](#). Attendance is generally opened 60 days before the next summit begins. The home page also includes links to the videos of sessions from previous summits. Enjoy the deep dives! —Andrew Binstock

Programming Devices

In this issue, our feature articles focus on how to program devices. The devices range from small, single-purpose hardware items to complex machines. When hooked together, typically via the cloud, they form the Internet of Things (IoT). Much of IoT programming consists of raw data processing: you gather data from the device, and you send it commands. The biggest challenges are accessing the device and knowing the command sequence that translates into specific actions. Learn how to do most of this, add a little back-end processing in the cloud, and you're most of the way there.


This issue has three articles that build on each other to show how this kind of programming is done in Java. The first article, by Gastón Hillar ([page 14](#)), shows the basics of accessing a device—in this case, an inexpensive Intel-based IoT board—and using its features by querying sensors and turning on various pins that light an LED. If you're not terribly familiar with Java access to hardware, this is the article to start with. It's an ideal hobbyist project: instructive and useful, and it can be completed in a single afternoon.

The second article ([page 24](#)) demonstrates how to control a CNC router by

sending commands to it from a Raspberry Pi board. CNC routers are devices that shape hardware items by grinding down or cutting a piece of material into the right shape. They're like 3-D printers except that they remove material from a block rather than iteratively build up the product from new material. As this article shows, once a connection to the device is established, the primary task is getting data from the device and sending it new commands. The Java code for managing these two activities is remarkably straightforward. The author, Stephen Chin, wrote a similar article in our May/June 2015 issue explaining how to interact with an electronic scale from a Raspberry Pi.

For the IoT to become a reality, devices need to share data and typically will do so by funneling the data from multiple devices to a cloud instance. Our third article on IoT ([page 32](#)) shows how this works on a cloud platform as a service custom-made for such data and for controlling devices remotely.

In addition, you'll find our regular assortment of Java articles: in-depth tutorials, explorations of JVM languages, book reviews, letters we've received, and of course, our Java quiz.





Capturing data and responding to it are the heart and soul of IoT. Here's how it works with Java and an inexpensive board that uses an x86 processor.

Prerequisites

For development, I use the Intel System Studio IoT Edition software. It is an Eclipse-based IDE that makes it easy to create a new IoT project with Java as the main programming language. I also use the latest available versions of both the mraa and upm libraries. The mraa library is a low-level skeleton library for communication on Linux platforms, and upm is a set of libraries for interacting with sensors and actu-

You can also discover the board and its services on the LAN automatically through the zero-configuration network-



A *photoresistor* is an electronic component also known as a light-dependent resistor (LDR) or photocell. It is not the best component for sensing ambient light with high levels of accuracy. However, the component is useful for determining

- A photoresistor.
- A $10,000\Omega$ (10 k Ω) resistor with 5 percent tolerance. The color bands for the resistor are brown, black, orange, and gold.
- A common cathode 5 mm RGB LED.

The class declares a method that translates a brightness level from 0 to 255 (inclusive) into the appropriate output duty-cycle value for the PWM pin.

Then, the constructor calls the `pwm.enable` method with `true` as a parameter to set the enable status of the PWM-enabled pin and allow the code to start setting the output duty-cycle percentage for the PWM pin with calls to the `pwm.write` method.

Finally, the constructor calls the `setBrightnessLevel` method with 0 as the value for the `brightnessLevel` argument. This way, the constructor sets the brightness level to 0 for the LED wired to the specified pin number and turns off the LED. Specifically,

the call turns off a specific color component of the RGB LED.

The class declares a `setBrightnessLevel` method that translates a brightness level from 0 to 255 (inclusive) into the appropriate output duty-cycle value for the PWM pin. The method receives a brightness level `int` value in the `brightnessLevel` argument.

First, the code makes sure that the brightness level is a value between 0 and 255 (inclusive). If the value is out of range, the code uses either the lower-level or the upper-level value and saves it in the `validBrightnessLevel` local variable.

Then, the code calculates the required output duty-cycle percentage for the PWM pin to represent the brightness level as a float value between 1.0f and 0.0f (100 percent and 0 percent). It is necessary to divide the valid brightness level (`validBrightnessLevel`) by 255f. The code saves the value in the `convertedLevel` variable. The next line calls the `this.pwm.write` method with the `convertedLevel` variable for the percentage argument and sets the output duty cycle for the pin configured as the PWM output to `convertedLevel`.

Finally, the code saves the valid brightness level in the `brightnessLevel` field, which is read-only accessible through the `getBrightnessLevel` method. The last line prints details about the brightness level set to the LED identified with a name and wired to a specific pin number. The line is printed with `System.out.format`, and it is possible to see the output when you run the generated JAR file through the IDE or by running commands through SSH on Yocto Linux running on the board. I'll dive deep on the benefits of printing useful information later.

Measuring Ambient Light via Analog Input

Now, I create a new `VoltageInput` class that represents a voltage source connected to an analog input pin on the board. The following lines show the code for this new class:

```
class VoltageInput {
    private final int analogPin;
    private final Aio aio;

    public VoltageInput(int analogPin) {
        this.analogPin = analogPin;
        this.aio = new Aio(analogPin);
        // Configure the ADC
        // (short for Analog-to-Digital Converter)
        // resolution to 12 bits (0 to 4095)
        this.aio.setBit(12);
    }

    public float getVoltage() {
        long rawValue = this.aio.read();
        float voltageValue =
            rawValue / 4095f * 5f;
        return voltageValue;
    }
}
```

```
public int getAnalogPin() {
    return analogPin;
}
```



```

        this.measureLight();
    }

    public void measureLight() {
        this.measuredVoltage =
            this.voltageInput.getVoltage();
        if (this.measuredVoltage <
            SimpleLightSensor.EXTREMELY_DARK) {
            this.lightLevel =
                SimpleLightSensor.LL_EXTREMELY_DARK;
        } else if (this.measuredVoltage <
            SimpleLightSensor.VERY_DARK) {
            this.lightLevel =
                SimpleLightSensor.LL_VERY_DARK;
        } else if (this.measuredVoltage <
            SimpleLightSensor.JUST_DARK) {
            this.lightLevel =
                SimpleLightSensor.LL_JUST_DARK;
        } else {
            this.lightLevel =
                SimpleLightSensor.LL_SUNNY_DAY;
        }
    }

    public String getLightLevel() {
        return this.lightLevel;
    }
}

```

The class defines the following three class fields, which specify the maximum voltage values that determine each light level.

- EXTREMELY_DARK: 2.1V
- VERY_DARK: 3.1V
- JUST_DARK: 4.05V

If the retrieved voltage is lower than 2.1V, the environment is extremely dark. If the retrieved voltage is lower than 3V,

the environment is very dark. If the retrieved voltage is lower than 4.05V, the environment is just barely dark. These values work with a specific photoresistor; you might have to check the voltage values that determine specific environments in your own configuration. It is necessary to change only the values of the class fields.

The main goal for the `SimpleLightSensor` class is to convert a quantitative value (a voltage value) into a qualitative value (an ambient light description). The class declares the following four class fields with the descriptions for the light levels:

- LL_EXTREMELY_DARK
- LL_VERY_DARK
- LL_JUST_DARK
- LL_SUNNY_DAY

When I create an instance of the `SimpleLightSensor` class, I need to specify in the `analogPin` argument, which is the analog pin number to which the voltage divider that includes the photoresistor is connected. The constructor creates a new `VoltageInput` instance with the received `analogPin` argument and saves its reference in the `voltageInput` field. The next line calls the `measureLight` method that transforms the voltage value retrieved with the `VoltageInput` instance (`this.voltageInput`) into a description of the light level.

The class declares the `measureLight` method, which saves the voltage value retrieved by calling the `this.voltageInput.getVoltage` method in the `measuredVoltage` field. Then, the next lines use the previously explained class fields to determine whether the value of the `measuredVoltage` field is lower than the maximum voltages that determine each light level. The code sets the appropriate value for the `lightLevel` field according to the measured voltage value. Then, it is

**The new class
allows me to
transform a voltage
value into a light
measurement and
description.**

possible to access the light level description by calling the `getLightLevel` method.

Controlling One Input and Three Outputs with a Board Manager

Now, I will create a new `BoardManager` class that creates an instance of the previously coded `SimpleLightSensor` class and three instances of the `VariableBrightnessLed` class. This way, there is one instance of the `VariableBrightnessLed` class for each color component of the RGB LED. The class fires actions when the ambient light changes. Specifically, the class adjusts the brightness for the three color components of the RGB LED based on the measured ambient light. The following lines show the code for this new class:

```
class BoardManager {
    public final SimpleLightSensor lightSensor;
    public final VariableBrightnessLed redLed;
    public final VariableBrightnessLed greenLed;
    public final VariableBrightnessLed blueLed;

    public BoardManager() {
        this.lightSensor =
            new SimpleLightSensor(0);
        this.redLed =
            new VariableBrightnessLed(6, "Red");
        this.greenLed =
            new VariableBrightnessLed(5, "Green");
        this.blueLed =
            new VariableBrightnessLed(3, "Blue");
    }

    public void setRGBBrightnessLevel(int value) {
        this.redLed.setBrightnessLevel(value);
        this.greenLed.setBrightnessLevel(value);
        this.blueLed.setBrightnessLevel(value);
    }
}
```


The `setRGBBrightnessLevel` method calls the `setBrightnessLevel` method for the three `VariableBrightnessLed` instances with the `value` received as an argument. This way, the three color components of the RGB LED are set to the same brightness level through a single call.

The `updateLedsBasedOnLight` method retrieves the light-level description from the `SimpleLightSensor` instance and calls the previously explained `setRGBBrightnessLevel` method to set the brightness level for the three components of the RGB LED based on the measured light. If it is extremely

dark, the brightness level is set to 255. If it is very dark, the brightness level is set to 128. If it is just dark, the brightness level is set to 64. Otherwise, the brightness level is set to 0, which means the RGB LED is completely off.

Now, I will write code that uses the `BoardManager` class to measure the

ambient light and set the brightness for the three color components of the RGB LED based on the measured ambient light. The following lines show the code for the new `AmbientLightAndLed` class:

```
public class AmbientLightAndLed {
    public static void main(String[] args) {
        String lastlightLevel = "";
        BoardManager board = new BoardManager();
        while (true) {
            board.lightSensor.measureLight();
            String newLightLevel =
                board.lightSensor.getLightLevel();
            if (newLightLevel != lastlightLevel) {
```

```
// The measured light level has changed
lastLightLevel = newLightLevel;
System.out.format(
    "Measured light level: %s\n",
    newLightLevel);
board.updateLedsBasedOnLight();
}
try {
    Thread.sleep(1000);
} catch (InterruptedException e) {
    System.err.format(
        "Sleep interruption: %s",
        e.toString());
}
}
}
```

The main goal for the SimpleLightSensor class is to convert a quantitative value (a voltage value) into a qualitative value (an ambient light description).

STEPHEN CHIN

Raspberry Pi-Controlled CNC Router

Programming the Raspberry Pi to manage cutting, carving, and routing

CNC routers are very useful tools for creating physical objects by carving material out of blocks of wood, plastic, or even soft metals. (*CNC* stands for *computer numerical control* and simply indicates that the routers are directed by computers.) They are the opposite of a 3-D printer, which creates objects by adding material, but they are equally useful. Besides cutting out flat objects, by the careful removal of material in small layers you can carve complex 3-D geometries into durable materials.

Traditionally CNC routing was limited to large industrial machines, but modern desktop CNC routers are small enough to fit beside your computer. They are relatively quiet and have full enclosures to limit the spread of dust and shavings. They use motors similar to 3-D printer motors, so they have very high accuracy that can be used to create fine detail in carvings. And they are relatively easy to control, because they accept standard G-code instructions from an attached computer or microcontroller. G-code is a simple text-based language for describing low-level machine instructions that control a CNC router, 3-D printer, or other machinery. There are G-code commands to control coordinates, movement, rotation, and other machine functions.

In this article, I show you how to interface with the Nomad 883 Pro router produced by Carbide 3D. This router accepts G-code over a serial connection and uses a Grbl controller on an embedded Arduino board. If you are using a different printer there will most likely be different G-code initialization instructions and possibly a different controller board, so check your printer specifications.

You can find the complete code for the example I use in this article on [GitHub](#).

Connecting to the Router

For sending data to the router, I use the UniversalGcode Sender project by Will Winder. It exposes a simple API for sending commands to an attached router that uses either a Grbl or TinyG controller.

To start, download the latest version of UniversalGcode Sender (version 1.0.9, as of the writing of this article) from the [download page](#).

Then create a new project in your favorite IDE (I happen to be using NetBeans) and include the UniversalGcodeSender.jar file as a dependency in your project. You can create a new connection to the router as follows:

```
static GrblController grblController;
static final CutterListener listener =
    new CutterListener();
static String PORT_NAME = "/dev/ttyACM0";

public static void main(String[] args) throws Exception {
    // Send decimals with "." instead of ",":

    Locale.setDefault(Locale.ENGLISH);
    grblController = new GrblController();
    grblController.addListener(listener);
    grblController.setSingleStepMode(true);
    Boolean openCommPort =
```

```

        grblController.openCommPort(PORT_NAME, 115200);
    if (openCommPort != true) {
        throw new IllegalStateException(
            "Cannot open connection to the cutter");
    }
}

```

`PORT_NAME` is a constant that specifies the serial port your router is using, and 115200 is the baud rate at which your router communicates. To determine the port name of your router on Mac OS X or Linux, check the output of the `dmesg` command right after connecting the router. On Windows, you can use Device Manager to determine the port name by checking the Ports section for attached devices.

Note that I am using single-step mode to avoid errors on the Nomad 883 Pro. It prefers commands to be sent sequentially rather than having them be queued.

Once you have determined the correct port name and initiated a connection to the router, the next step is to initiate the router's homing and tool-measurement step. To initialize the Nomad 883 Pro, I execute the following G-code commands:

```
static final List<String> PROBE1 =
    Arrays.asList("G4P0.005", "M05", "G92.1",
        "G54", "G10 L2 P1 X0 Y0 Z0", "G21",
        "G49", "G90", "G10 L2 P1 X0 Y0 Z0",
        "G0 X-2.5 Z-5", "G0 Z-35.000",
        "G38.2Z-105 F800", "G4P0.005");

static final List<String> PROBE2 =
    Arrays.asList("G0 Z-70",
        "G38.2Z-182.675F200.0", "G4P0.005");

static final List<String> PROBE3 =
    Arrays.asList("G0 Z-5", "G0 X-5");
```

PROBE1 initializes the coordinate system and moves to the probe location. Then it checks the probe length at a rate of 800 mm/m.

PROBE2 repeats the

probing at a slower, more accurate speed of 200 mm/m. Finally,

PROBE3 moves the head away from the work surface and back to the home position.

Because the `Grb1Controller` is designed to be asynchronous, issuing commands requires a little bit of work. To simplify this process, I created some wrapper commands that send commands and wait for them to finish. Using those wrapper methods, here is the initialization cycle:

```
waitForConnection();
homeAndWait();
sendSequenceAndWait(PROBE1);
sendSequenceAndWait(PROBE2);
sendSequenceAndWait(PROBE3);
```

Before you can issue commands, the initialization cycle of the router has to finish. For this, I wait until I get back the initialization messages from the CNC router:

```
private static void waitForConnection()
    throws InterruptedException {
    synchronized (listener) {
        while (!listener.connected) {
            listener.wait();
        }
    }
}
```

Create a new project in your favorite IDE (I happen to be using NetBeans), and include the UniversalGcodeSender.jar file as a dependency in your project.


```

        prbZ =
            Double.parseDouble(matcher.group(1));
    }
}

// Additional empty methods omitted
}

```

Executing this program homes the router and runs the probe algorithm to test the length of the tool. The last method parses the probe output to find the Z coordinates of the router when the end-mill's tip is touching the probe. To output the probe coordinates and cleanly shut down the router, you can finish by using the following code:

```
System.out.println(
    "PRB_Z = " + listener.prbZ);
TimeUnit.SECONDS.sleep(5);
grblController.closeCommPort();
```

Write down the value of `PRB_Z`, because you will need it for calculating the offset to the workplane.

Calculating the Workspace Location

To figure out the workspace location, you need to manually move the router to touch the surface of the work area. The easiest way to do this is to open the [UniversalCodeSender](#) user interface. To run it, use the appropriate platform-specific script or simply type the following at the command line:

```
java -jar UniversalGcodeSender.jar
```

The UI lets you specify the port and the baud rate. Choose the same values you used in your Java program to connect to the

CNC router, and then click the **Open** button.

Once the router is connected, it is in a “WARN” state, because the router has not been homed. Navigate to the **Machine Control** tab and click the **\$H** button to send a homing command.

Now you can manually control the router head by using the **X**, **Y**, and **Z** movement buttons. Start with a course movement speed to get the head close to the work surface, and then change to a finer movement speed as you get closer. Using a piece of paper, get the head as close to the work surface as you can without touching the paper, as shown in **Figure 1**.

The value for Z shown under the **Machine Position** section of the screen tells you where the top of your work surface is. Record this value as **WRK Z**.

To calculate the **PROBE OFFSET**, use the following formulas:

$$\text{PROBE OFFSET} = \text{PRB Z} - \text{WRK Z} - 5$$

Plugging in the values that I got from my machine gives me:




Figure 1. Putting the head close to the work surface

ends the program (M30).

I also need to support multiple passes through the material, because acrylic is too hard to cut in one pass. For a 1/8-inch thick piece of acrylic, I recommend seven passes, each incrementally deeper. The following code calls the `drawStar` function multiple times to accomplish that:

```
sendSequenceAndWait(START_SPINDLE);
moveToStart(50, 100);
for (int i = 1; i <= Z_STEPS; i++) {
    double newZ = MATERIAL_THICKNESS *
        (Z_STEPS - i) / Z_STEPS;
    sendCommandAndWait(
        "G1Z" + String.format("%.3f", newZ) +
        "F355.600");
    sendCommandAndWait("F1117.600");
    sendSequenceAndWait(
        drawStar(9, 50, 90, 100));
}
sendSequenceAndWait(END_SEQUENCE);
```

This code makes use of a new argument, `F`, to `G1`. It is the feed rate. You can also call this code by itself to set the feed rate for subsequent linear operations. Finally, there are two parameters, the `MATERIAL_THICKNESS` and number of `Z_STEPS`, that you can set based on the material you are using:

```
static final int Z_STEPS = 7;
static final double
    MATERIAL_THICKNESS = 25.4 * 1/8;
```

Just to make sure the program is working correctly, you might want to perform a first run with no material inside the machine. If anything unexpected happens, killing the Java program or pressing the power switch on the Nomad 883 Pro will immediately stop the operation.

Once you are confident that you have a working application, load a piece of acrylic plastic into your Nomad 883 Pro, and run the program again. The router should slowly cut the acrylic layer by layer until it gets just above the router's wasteboard.

Upon completion, vacuum up the plastic scraps and carefully remove the finished piece from the tray. You should have a perfectly cut star similar to the one shown in **Figure 2**.

Running the Code on the Raspberry Pi

Running the code on the Raspberry Pi is the easiest part of the process described in this article. Java 8 comes bundled with the standard Raspbian Linux distribution, so if you have an up-to-date installation, you already have Java ready to go. And if you are using NetBeans, there is built-in support for running the code on the Raspberry Pi using the Remote Java Platform functionality.

To set up your Raspberry Pi in NetBeans for the first time:

1. In the **Tools** menu, choose **Java Platforms**.
2. Click **Add Platform** and select **Remote Java Standard Edition**.




Figure 2. An acrylic star cut using the presented code



By pushing device control and analytics into an IoT-specific cloud, devices can be remotely managed and operated.

- The cloud is more than someone else's data center. Look past infrastructure and toward platform as a service (PaaS), where you're insulated from details such as servers, load balancers, clusters, database sizes, and configuration. Instead, focus on solving problems and writing code.
- PaaS services offer ways to implement analytics that don't replace the fun of Java coding, but instead reduce tedious tasks.
- The cloud puts the power of customer- and business-driven changes in the hands of business owners. This reduces the need for developers to make minor changes and tweaks. For instance, by offloading key analytics from within Java to the cloud, you leave the task of tuning threshold values, monitoring details, and dealing with other tunable param-

Looking at the IoTCS [documentation](#), you'll notice many REST APIs, rules-based analytics, and integration with business intelligence, mobile devices, and enterprise systems based on packages and adapters. It might not be apparent how Java plays a role. However, Oracle provides a Java-based IoTCS [Client Software Library](#) that facilitates device, gateway, and enterprise application development in Java.



For example, IoTCS uses OAuth 2.0 tokening to implement much of its security model. Using the REST APIs, your device would need to read an SSL keystore, Base64-encode the appropriate username and password, send the proper OAuth 2 handshake messages, and handle the responses—all to activate an IoT device and have it communicate data. Doing this with pure REST can be tedious. However, by using parts of the Oracle IoT Cloud Service Java Client Software Library, the code to authenticate and activate an IoT device is reduced to what's shown in Listing 1.

■ Listing 1.

```
String url = ...
HttpClient httpClient =
    new HttpClient(server_url,uid,pwd);

WebResource resource =
    httpClient.createHttpClient(true).resource(url);

TokenInputDetails tokenInputDetails =
    new TokenInputDetails();

tokenInputDetails.setTokenType(
    TokenType.TOKEN_TYPE_ACTIVATION);
tokenInputDetails.setDeviceId(endpointID);
tokenInputDetails.setSharedSecret(secret);

String deviceModelStr =
    "urn:com:acme:conveyorbelmodel";
```

Oracle provides a Java-based Oracle IoT Cloud Service Client Software Library that facilitates device, gateway, and enterprise application development in Java.

```
String messageFormat =  
    "urn:com:acme:conveyorbeltnode:speed";  
  
Authorization authorization =  
    new Authorization(resource);  
  
PrivateKey privateKey =  
    authorization.activate(tokenInputDetails);
```

In fact, this code can be reduced further to the following using the `oracle.iot.client.device.DirectlyConnectedDevice` class from the IoTCS Client Software Library (see **Listing 2**).

■ **Listing 2.**

```
DirectlyConnectedDevice device =
    new DirectlyConnectedDevice();
if ( ! device.isActivated() ) {
    device.activate( getDeviceModelURN() );
}
```

When leveraging the cloud-based development paradigm, the key is understanding which aspects of an IoT solution to leave as cloud-handled analytics and which to implement in code. The goal is to alleviate tedious maintenance requirements by allowing the business to make changes without new software deployments. To demonstrate, I describe an IoT application scenario, dive into the Java code, and explore the Oracle IoT Cloud Service setup that makes this goal achievable.

Sample Application: Monitoring Industrial Conveyor Belts

In this sample IoT application, a factory conveyor belt is implemented, monitored, and controlled using Java and IoTCS. Although the belt motor is emulated here, the Java code is written with the IoTCS Client Software Library to run on an actual device or on a gateway connecting the device to the cloud. It's assumed that the conveyor belt in this scenario has the ability to connect to the internet directly to send

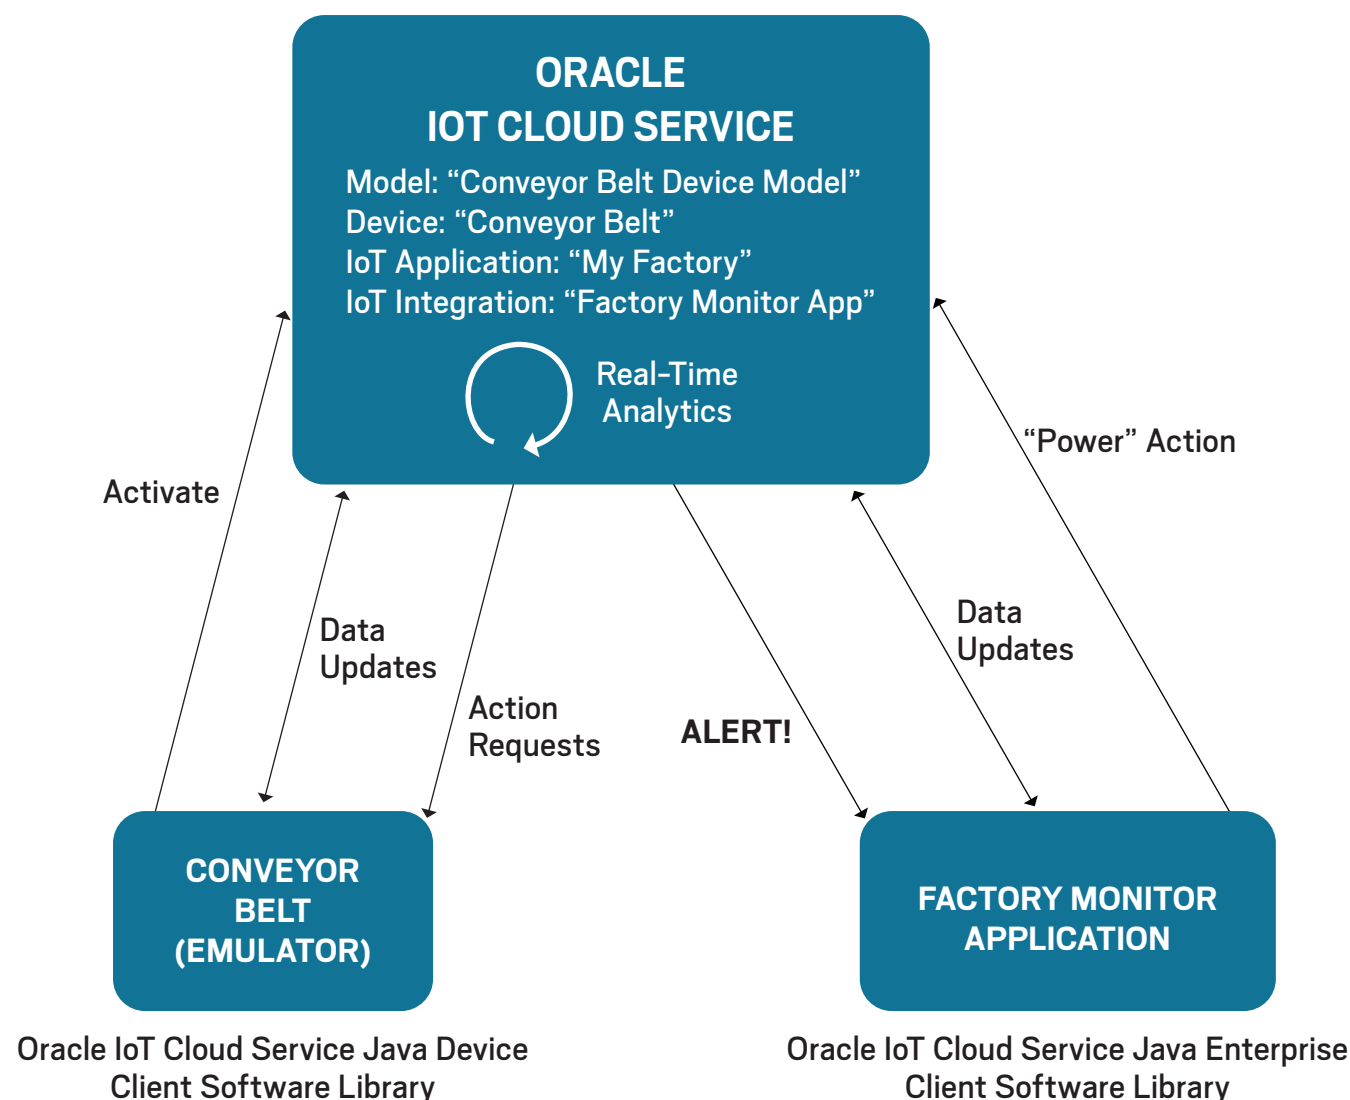


Figure 2. The architecture of the sample Oracle IoTCS application




Figure 3. A Streams Explorer analytic definition

with the appropriate device attributes.

The data source is the conveyor belt motor, and the logic should send an alert when the motor temperature goes above 100 degrees Fahrenheit and the cooling fan speed is less than 1 (the fan is not on). The TempThreshold2 analytic is similar, except that it looks for a higher temperature (> 150 degrees Fahrenheit) and a fan speed that is less than 2. The remaining analytics are similar.

However, these analytics are useless unless there's an application that listens for the resulting alerts to act on them. This is the Java monitoring application mentioned previously, which is integrated through a REST endpoint to which IoTCS sends the analytics-driven alerts. As a result, the monitoring application instructs the device to change its belt motor speed, adjust the cooling fan speed, or shut down in response to the analytics alerts (see **Listing 3**).

■ **Listing 3.**

```
public static void handleREST(HttpExchange t)
    throws Exception {
    // Read request JSON string
    BufferedReader bodyReader =
        new BufferedReader(
            new InputStreamReader(
                t.getRequestBody() ));
    //...

    // Parse the JSON string
    StringReader sr =
        new StringReader(alert.toString());
    JsonReader jr = Json.createReader(sr);
    JsonObject jo = //...
    JsonObject payload =
        jo.getJsonObject("payload");
```

```
// Get the alert type by URN
```



```
String alertType =
    payload.getString("format");

switch ( alertType ) {
    case TEMP_TOO_LOW_ALERT_URN:
        device.set(FAN_SPEED_ATTR,
                   OFF);
        break;
    case TEMP_LOW_ALERT_URN:
        device.set(FAN_SPEED_ATTR,
                   FAN_LOW_SPEED);
        break;
    case TEMP_MED_ALERT_URN:
        device.set(FAN_SPEED_ATTR,
                   FAN_MED_SPEED);
        break;
    case TEMP_HIGH_ALERT_URN:
        device.set(FAN_SPEED_ATTR,
                   FAN_HIGH_SPEED);
        break;
    case TEMP_CRITICAL_ALERT_URN:
        shutDownBelt(device);
        break;
    //...
}
```

With this IoTCS analytics-based implementation, changing the thresholds that indicate when the cooling fan speed needs to be increased or decreased doesn't require changes to the Java application. The analytics can be set and adjusted instead by someone who understands the mechanical properties of the belt motor. The alternative would be the much less satisfactory hard-coding of values and limits, using code such as shown in **Listing 4**.

■ **Listing 4.**

```
VirtualDevice device = event.getVirtualDevice();
int temp = (Integer)namedValue.getValue();
int coolingFanSpeed = device.get(FAN_SPEED_ATTR);

// Determine if action is needed based on new temp
if ( temp >= 300 ) {
    // Shut things down!
    device.set(BELT_SPEED_ATTR, 0);
}
else if ( temp > 250 && ...
```

To receive notification that the monitoring application changed the belt or fan speed, the device needs to register for appropriate attribute changes, as shown in **Listing 5**.

■ Listing 5.

```
DeviceModel model =
    device.getDeviceModel(
        CONVEYOR_BELT_MODEL_URN);

VirtualDevice virtualDevice =
    device.createVirtualDevice(
        device.getEndpointId(),
        model);

virtualDevice.setOnChange(
    BELT_SPEED_ATTR,
    new VirtualDevice
        .ChangeCallback<VirtualDevice>() {
            public void onChange(
                VirtualDevice
                    .ChangeEvent<VirtualDevice> e) {
                onBeltSpeedChange(e);
            }
        });
```


IoTCS analytics alert the monitoring application regarding when to turn the cooling fan on or off, as well as about other important events such as low oil pressure and overheating

Eric Bruno is a principal sales consultant and an Oracle IoT Champion in the Oracle ISV software sales team. He has more than 20 years' experience in the information technology community as an enterprise architect, developer, and industry analyst with expertise in full lifecycle, large-scale software architecture, design, and development.

[Oracle IoT Cloud Service home page](#)

OAuth 2.0



A look at what's coming next in JAX-RS 2.1

SSEs are a new technology that is part of HTML5. SSEs automatically receive updates from a server via HTTP, and they offer an easy-to-use, one-way streaming communication

The server side of SSEs. The following code shows an SSE-enabled JAX-RS endpoint.



The `SseEventSource` is created by calling the `target` method on a `WebTarget`, registering an `SseEventSource.Listener`, and opening the connection. After successfully opening the connection, the current thread continues and the listener—which in this case would be `dataConsumer.accept`—is called as soon as events arrive.

`SseEventSource` handles all required plumbing, including reconnecting after a connection loss, by sending an appropriate Last-Event-ID header and then handling Retry-After headers sent from the server appropriately.

If a more sophisticated way is needed, for instance, to manually control the Last-Event-ID header, you can manually request an `SseEventInput` from the server.

```
public class StatefulSseClient {

    private final WebTarget target =
        ClientBuilder.newClient().target("...");
    private final Consumer<String> dataConsumer;
    private String lastEventId;
    private SseEventInput eventInput;

    public StatefulSseClient(
        Consumer<String> dataConsumer) {
        this.dataConsumer = dataConsumer;
    }

    public void start() {
        eventInput = target
            .request(MediaType.SERVER_SENT_EVENTS)
```

```

        .header(HttpHeaders.LAST_EVENT_ID_HEADER,
                lastEventId)
        .get(SseEventInput.class);

    new Thread(() -> {
        while (!eventInput.isClosed()) {
            final InboundSseEvent event =
                eventInput.read();
            if (event != null) {
                lastEventId = event.getId();
                dataConsumer.accept(
                    event.readData());
            }
        }
    }).start();
}

public void stop() {
    if (eventInput != null &&
        !eventInput.isClosed())
        try {
            eventInput.close();
        } catch (IOException e) {
            // suppress
        }
}

```

By requesting the input directly, all required information has to be sent manually and potential reconnects need to be handled manually. Therefore, the client sends the text/event-stream HTTP header and expects the response to be an `SseEventInput` type that is handled appropriately by the JAX-RS implementation. This event input is used to receive the actual incoming events.

46



MICHAEL KÖLLING

The Evolving Nature of Interfaces

Understanding multiple inheritance in Java

In the “New to Java” series, I try to provide benefit by picking topics that invite a deeper understanding of the conceptual background of a language construct. Often, novice programmers have a working knowledge of a concept—that is, they can use it in many situations, but they lack a deeper understanding of the underlying principles that would lead to writing better code, creating better structures, and making better decisions about when to use a given construct. Java *interfaces* are just such a topic.

In this article, I assume that you have a basic understanding of inheritance. Java interfaces are closely related to inheritance, as are the `extends` and `implements` keywords. So, I will discuss why Java has two different inheritance mechanisms (indicated by these keywords), how abstract classes fit in, and what various tasks interfaces can be used for.

As is so often the case, the story of these features starts with some quite simple and elegant ideas that lead to the definition of concepts in early Java versions, and the story gets more complicated as Java advances to tackle more-intricate, real-world problems. This leads to the introduction of default methods in Java 8, which muddy the waters a bit.

A Little Background on Inheritance

Inheritance is quite straightforward to understand in principle: a class can be specified as an extension of another class. In such a case, the present class is called a *subclass*, and the class it’s extending is called the *superclass*. Objects of the subclass have all the properties of both the superclass and the subclass. They have all fields defined in either subclass or

superclass and also all methods from both. So far, so good.

Inheritance is, however, the equivalent of the Swiss Army knife in programming: it can be used to achieve some very diverse goals. I can use inheritance to reuse some code I have written before, I can use it for subtyping and dynamic dispatch, I can use it to separate specification from implementation, I can use it to specify a contract between different parts of a system, and I can use it for a variety of other tasks. These are all important, but very different, ideas. It is necessary to understand these differences to get a good feel for inheritance and interfaces.

Type Inheritance Versus Code Inheritance

Two main capabilities that inheritance provides are the ability to inherit code and the ability to inherit a type. It is useful to separate these two ideas conceptually, especially because standard Java inheritance mixes them together. In Java, every class I define also defines a type: as soon as I have a class, I can create variables of that type, for example.

When I create a subclass (using the `extends` keyword), the subclass inherits both the code and the type of the superclass. Inherited methods are available to be called (I’ll refer to this as “the code”), and objects of the subclass can be used in places where objects of the superclass are expected (thus, the subclass creates a subtype).

Let’s look at an example. If `Student` is a subclass of `Person`, then objects of class `Student` have the type `Student`, but they also have the type `Person`. A student is a person. Both the code and the type are inherited.




Figure 1. An example of multiple inheritance




Figure 2. An example of diamond inheritance

The decision to link type inheritance and code inheritance in Java is a language design choice: it was done because it is often useful, but it is not the only way a language can be designed. Other programming languages allow inheriting code without inheriting the type (such as C++ *private inheritance*) or inheriting type without code (which Java also supports, as I explain shortly).

Multiple Inheritance

The next idea entering the mix is *multiple inheritance*: a class may have more than one superclass. Let me give you an example: PhD students at my university also work as instructors. In that sense, they are like faculty (they are instructors for a class, have a room number, a payroll num-

ber, and so on). But they are also students: they are enrolled in a course, have a student ID number, and so on. I can model this as multiple inheritance (see **Figure 1**).

`PhDStudent` is a subclass of both `Faculty` and `Student`. This way, a PhD student will have the attributes of both students and faculty. Conceptually this is straightforward. In practice, however, the language becomes more complicated if it allows multiple inheritance, because that introduces new problems: What if both superclasses have fields with the same name? What if they have methods with the same signature

but different implementations? For these cases, I need language constructs that specify some solution to the problem of ambiguity and name overloading. However, it gets worse.

Diamond Inheritance

A more complicated scenario is known as *diamond inheritance* (see Figure 2). This is where a class (`PhDStudent`) has two superclasses (`Faculty` and `Student`), which in turn have a common superclass (`Person`). The inheritance graph forms a diamond shape.

Now, consider this question: if there is a field in the top-level superclass (`Person`, in this case), should the class at the bottom (`PhDStudent`) have one copy of this field or two? It inherits this field twice, after all, once via each of its inheritance branches.

The answer is: it depends. If the field in question is, say, an ID number, maybe a PhD student should have two: a student ID and a faculty/payroll ID that might be a different number. If the field is, however, the person's family name, then you want only one (the PhD student has only one family name, even though it is inherited from both superclasses).

In short, things can become very messy. Languages that allow full, multiple inheritance need to have rules and constructs to deal with all these situations, and these rules are complicated.

Type Inheritance to the Rescue

When you think about these problems carefully, you realize that all the problems with multiple inheritance are related to inheriting code: method implementations and fields. Multiple code inheritance is messy, but multiple type inheritance causes no problems. This fact is coupled with another observation: multiple code inheritance is not terribly important, because you can use *delegation* (using a reference to another object) instead, but multiple subtyping is often very useful and not easily replaced in an elegant way.

Col 0	Col 1	Col 2	Col 3	Col 4
0 x 0	0 x 1	0 x 2	0 x 3	0 x 4
1 x 0	1 x 1	1 x 2	1 x 3	1 x 4
2 x 0	2 x 1	2 x 2	2 x 3	2 x 4

Figure 1. A sample table created with Fantom domkit

just like Swing. The domkit [Table](#) automatically handles all the thorny issues: efficiently mapping the model to DOM elements (but only for the visible rows), scrolling, column sorting, single/multiple selection, and many other features.

Let's look at some real code to illustrate how to build a table using domkit:

```
@Js
class MyTableModel : TableModel
{
    override Int numCols() { 5 }
    override Int numRows() { 3 }
    override Void onHeader(Elem e, Int c) {
        e.text = "Col $c" }
    override Void onCell(
        Elem e, Int c, Int r, TableFlags f)
        { e.text = "$r x $c" }
}
```

This creates a subclass of the domkit `TableModel`. It defines the number of rows and columns for the model and provides a callback for how to render the column headers and cells.

Now let's see how I put it all together to create a table:

```
table := Table
{
    model = MyTableModel()
    sel.multi = true
    onAction |t| {
```

```
        echo("onAction: ${t.sel.indexes}") }
onSelect |t| {
    echo("onSelect: ${t.sel.indexes}") }
}
```

This code creates an instance of `Table` using the `model` class and sets multiple selection to be enabled. Then, I add some event handlers for action (double click) and selection changes that echo to stdout the selected row indexes. **Figure 1** shows what the table looks like in a browser.

If you have experience building HTML5 UIs and miss the higher-level abstractions that a widget toolkit such as Swing provides, then domkit might be just the technology for you.

Conclusion


In this article, I examined four key features of Fantom: portability, immutability, actor concurrency, and domkit. This only briefly touches on the Fantom language, libraries, and tools. If you would like to learn more, visit our [website](#), where you can find documentation, a community forum, an active mailing list, and links to downloads as well as our BitBucket repo. `</article>`

Brian Frank is the founder and president of SkyFoundry, a software company specializing in IoT data collection, analysis, and visualization. Brian and his brother, Andy Frank, have been developing the Fantom platform since 2005. Brian also serves as the technical lead for project-haystack.org, an open source project for defining data models and formats in the IoT space.

learn more

Why Fantom?

Written by leading Java experts, Oracle Press books offer the most definitive, complete, and up-to-date coverage of Java available.



Raspberry Pi with Java: Programming the Internet of Things (IoT)


Stephen Chin, James Weaver

Use Raspberry Pi with Java to create innovative devices that power the internet of things.




Introducing JavaFX 8 Programming **Herbert Schildt**

Learn how to develop dynamic JavaFX GUI applications quickly and easily.



Java: The Complete Reference, Ninth Edition **Herbert Schildt**

Fully updated for Java SE 8, this definitive guide explains how to develop, compile, debug, and run Java programs.



OCA Java SE 8 Programmer I Study Guide (Exam 1Z0-808) **Edward Finegan, Robert Liguori**

Get complete coverage of all objectives for Exam 1Z0-808. Electronic practice exam questions are included.

More subtle questions from an author of the Java certification tests

I've put together more problems that simulate questions from the [1Z0-809 Programmer II exam](#), which is the certification test for developers who have been certified at a basic level of Java 8 programming knowledge and now are looking to demonstrate more-advanced expertise. [Readers wishing basic instruction should consult the "New to Java" column, which appears in every issue. —Ed.] These questions might require careful deduction to obtain the right answer.

```
public IntSupplier doStuff(int [] vals, int i) {
    // line n1
    return () -> vals[i];
}
```

- a. The code compiles.
- b. The argument list must be changed to `(final int[] vals, final int i)` to allow the code to compile.
- c. If the code `if (vals[0] < 0) vals[i] = 0;` is added at line n1, the argument list must be changed to `(final int[] vals, final int i)` to allow the code to compile.
- d. The code `if (vals[0] < 0) vals[i] = 0;` can be added at line n1 without causing compilation errors.
- e. The code `vals = Arrays.copyOf(vals, vals.length);` can be added at line n1 without causing compilation errors.

```
public class Wrapper {
    public class Wrapped {}
}
```

- a. An instance of `Wrapped` can be created only by code inside the class `Wrapper`.
- b. An instance of `Wrapped` can be created using the expression `new Wrapper.Wrapped()`.
- c. An instance of `Wrapped` can be created using the expression `new Wrapper().Wrapped()`.
- d. An instance of `Wrapped` can be created using the expression `new Wrapper().new Wrapped()`.
- e. An instance of `Wrapped` can be created using the expression `new Wrapper::Wrapped()`.

```
public static int sumAndPrint(IntStream is) {
    int total = 0;
    is.parallel()
        .peek(v -> total += v)
        .forEach(System.out::println);
    return total;
}
```


Which is true? Choose one.

- a. The code reliably returns the correct sum of the numbers as it is.
- b. The code would reliably return the correct sum of the numbers if the body of the method were changed to this:

```
int[] total = {0};  
is.peek(v -> total[0] += v)  
    .forEach(System.out::println);  
return total[0];
```
- c. The code would reliably return the correct sum of the numbers if the type of `total` were changed to `Integer`.
- d. The code would reliably return the correct sum of the numbers if the call to `.parallel()` were moved after the call to `.peek(v -> total += v)`.
- e. The code would reliably return the correct sum of the numbers if the type of `total` were changed to `LongAdder` and if the lambda in `peek` were changed accordingly.


Question 4. Given this code:

```
public static void main(String[] args) {
    ExecutorService es = Executors.newFixedThreadPool(2);
    Callable<String> job = () -> {
        Thread.sleep(5000); // line n1
        return "Hello";
    };
    es.submit(job);
    es.submit(job);
    Future<String> handle = es.submit(job); // line n2
    System.out.println("submitted");
    String message = handle.get(); // line n3
    System.out.println("> " + message);
    System.exit(0);
}
```

Which is true? Choose one.

- a. Compilation fails due to an error at line n1.
- b. Compilation fails due to an error at line n3.
- c. The program throws a `RejectedExecutionException` because there are no available threads at line n2.
- d. The program immediately prints `submitted`, and then after a 10-second pause, it prints `> Hello`.
- e. The program pauses five seconds, then prints `submitted`, and then after a further five-second pause, it prints `> Hello`.

Question 5. Given that the current working directory of the following program contains this tree:



where one, two, a, b, and c are directories, and the program code is this:

```
public static void main(String[] args) throws Throwable {
    Files
        .find(Paths.get("."), 1, (p,a)->a.isDirectory())
        .forEach(System.out::println);
}
```

What is the output? Choose one.

- a. -
b. .
./one
./two

use of method references anyway. Further, the `new` keyword should follow the double colon in a method reference. Option E is just wrong.

Question 3. The correct answer is option E. This question investigates the concurrent mode of stream processing and the rules that ensure it runs correctly and in a thread-safe fashion. A common recommendation is that stream processing, if there's any chance of the stream being executed in concurrent mode, should not mutate any shared data. Such an approach is a great way to ensure that concurrency problems do not arise. However, the actual rule is somewhat less restrictive. If shared state will be modified, then the programmer must ensure that those modifications are safe. In this case, using a `LongAdder` would achieve this goal, and option E is the correct answer. Interestingly, this situation could also be served by using the `AtomicInteger`, but the `LongAdder` was deliberately created to support many concurrent mutations in a scenario where relatively few reads are performed. Both would be functionally correct, but the adder will be more scalable.

Let's look at the wrong answers and see what might be interesting about those. First, why doesn't this code compile? The variable `total` is a method-local variable, and for that to be used in an enclosed lambda, it must be `final` or effectively `final`. Consequently, because it is not `final`, and it cannot be because it is mutated, the code as it stands will not compile. This is why option A is incorrect.

Option B looks tempting, and in many cases it would probably work. First, it uses an array of one `int` to accumulate the total. This successfully sidesteps the problem of a final variable, because the variable is a pointer to the array and is, indeed, now effectively final. However, even though the call to `parallel` was also removed from the code, the stream was received as an argument to the method, so it's not safe to assume that the stream is now running sequentially. It could have been set to `parallel` by the caller. In that situation,

the code remains unsafe from a concurrency perspective. So, option B is incorrect, because the behavior would still not be reliable.

Option C is unworkable; `Integer` objects are immutable, and given that `total` must be effectively final, there's really no way that such a suggestion could result in correct counting.

Option D might seem tempting. Presumably the idea is that if the mutation operation, performed by `peek`, can be performed sequentially, there will be no concurrency issue with respect to the updates to `total`. Unfortunately, this fails for two reasons: first and most compellingly, the option still doesn't compile because `total` is (still) not effectively final. Second, and also important, is the fact that streams don't shift between parallel and serial modes along their length. The whole stream, from end to end, is either parallel or sequential. Therefore, moving the call to `parallel` later in the chain changes absolutely nothing. For both these reasons, option D is also incorrect.

Question 4. The correct answer is option B. There's quite a lot of code in this question. That is unusual in the actual certification exam, but it does happen. This question tests one specific piece of understanding, but it introduces a number of interesting side issues for discussion as distractions.

The code does not compile, simply because the `get` method on a `Future` object throws checked exceptions. In fact, the documentation declares three exceptions, two of which are checked. The unchecked exception is `CancellationException`, which indicates that the job was canceled. Therefore, trying to get its result is meaningless. The checked exceptions are `InterruptedException` and `ExecutionException`. An `InterruptedException` results if the `get` method, which will block if the job hasn't finished yet, is interrupted while waiting. An `ExecutionException` arises if the job itself throws an exception; the cause of the `ExecutionException` is the actual exception thrown by the job.

