

Dyalog APL

Experimental Functionality

Parallel Language Features

Version 14.0

Dyalog Limited

Minchens Court, Minchens Lane
Bramley, Hampshire
RG26 5BH
United Kingdom

tel: +44(0)1256 830030

fax: +44 (0)1256 830031

email: support@dyalog.com

<http://www.dyalog.com>

Dyalog is a trademark of Dyalog Limited
Copyright © 1982-2014



*Dyalog is a trademark of Dyalog Limited
Copyright © 1982 - 2014 by Dyalog Limited.
All rights reserved.*

*Isolate version 1.0.851
Dyalog Version 14.0*

Revision: 20140913_140

No part of this publication may be reproduced in any form by any means without the prior written permission of Dyalog Limited, Minchens Court, Minchens Lane, Bramley, Hampshire, RG26 5BH, United Kingdom.

Dyalog Limited makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Dyalog Limited reserves the right to revise this publication without notification.

UNIX is a registered trademark of The Open Group.

All other trademarks and copyrights are acknowledged.

Contents

1	INTRODUCTION	1
2	PHILOSOPHICAL PREAMBLE	2
2.1	APL is a Parallel Notation	2
2.2	Automatic Parallelisation	2
2.3	Explicit Parallelisation	3
3	INTRODUCING FUTURES AND ISOLATES	4
3.1	Isolates	4
3.2	Futures	5
3.3	Parallel Operator	6
4	INTRODUCING THE V14.0 IMPLEMENTATION	7
4.1	Disable "Pause on Error"	7
4.2	Parallel Each	8
4.3	Explicit or Persistent Isolates	9
4.4	Blocking	10
4.5	Errors in Asynchronous Expressions	10
4.6	Tracking the Status of Asynchronous Expressions	12
4.7	Interrupting Asynchronous Expressions	12
4.8	Debugging Asynchronous Expressions	13
5	CALLING BACK FROM ISOLATES	15
5.1	Calling Back to the Parent	15
5.2	Call-backs versus Interrupts	16
5.3	Calling From One Isolate to Another	16
6	ISOLATE PROCESSES	17
6.1	isolate.State	17
6.2	Limiting the number of Isolates per Process	17
6.3	Changing the number of Processes	18
6.4	Running Isolate Servers	18
6.5	Using an Isolate Server	19
6.6	Combining Local and Remote Isolates	19
7	DIFFERENCES BETWEEN NAMESPACES AND ISOLATES	21
7.1	File Ties are not shared with Isolates	21
7.2	Namespace "Refs" and Isolates	21
7.3	Classes versus Isolates	22
7.4	Unexpected output of Shy Results	22
8	OPTION SETTINGS	23
8.1	drc	24
8.2	homeport and homeportmax	24
8.3	isolates	24
8.4	listen	24

8.5	maxws	24
8.6	onerror	24
8.7	processors	25
8.8	processes.....	25
8.9	workspace	25
8.10	workspace	25
9	FURTHER READING	26
9.1	Tools and Samples on GitHub	26
9.2	Topics yet to come	26
APPENDIX A: FUNCTION AND OPERATOR REFERENCE		27
A.1	Models of Future Primitives.....	27
	isolate←∅ init	r←⊗ isolate.New ω..... 27
	The result is a reference to the isolate which has been created.	27
	f ll	fll isolate.ll..... 27
	f llX	fll" isolate.llEach..... 27
	f o ll	o.(fll) isolate.llOuter..... 27
	f llö ranks.....	fllö ranks isolate.llRank..... 28
	f llD	fll □ isolate.llKey..... 28
A.2	Isolate Server Management.....	28
	r←AddServer address ports.....	28
	r←RemoveServer address	29
	r←StartServer dummy.....	29
A.3	Configuration and Process State	30
	r←Config property [value].....	30
	r←Reset 0	30
	r←State "	30
A.4	State of Futures.....	32
	r←Available name	32
	r←Failed name	32
	r←Running name.....	32
	r←[null] Value name.....	32
APPENDIX B: TROUBLESHOOTING		33
B.1	"Normal" Errors.....	33
B.2	Initialisation and Communication Failures.....	33
B.3	Internal Errors	34

1 Introduction

This document describes proposals for two new primitives and a new type of array, designed to allow APL developers to take advantage of using multiple processors to execute code in parallel:

- The function “isolate” (\boxtimes) creates a special kind of namespace known as an *isolate* inside which all APL language statements are executed in parallel with the main APL process.
- The monadic operator “parallel” (\parallel) derives a function that will execute the left operand in an empty isolate.
- A new type of array known as a “future” is a pointer to the future result of an expression currently being executed in an isolate. Futures can be passed around as function arguments and as items of nested arrays, but will “block” any expression which actually needs the value pointed to by the future.

In version 14.0, *futures* are implemented as a new primitive array element type. However, the *isolate* function and the *parallel* operator are provided as models which are partly implemented in APL, provided in the distributed workspace `ws\i s o l a t e`.

Important: Although version 14.0 only provides a model, Dyalog Ltd. believes that the design is very close to being final, and urges all users with applications containing potential parallelism to make use of the experimental implementation, to help confirm the design and reveal any weaknesses that need to be worked on in the final implementation as part of the core language.

The current plan (subject to change without notice) is to produce a revised model based on user feedback for version 14.1, and a fully “primitive” implementation in version 15.0.

Dyalog Ltd. is aware that attempting to use the Workspace Explorer or attempting to edit isolate contents may cause the interpreter to terminate unexpectedly. A fix for this will be issued in due course.

1.1 isolate workspace version 1.0.851, September 13 2014

The first update to the isolate workspace has been released in September 2014, with the following enhancements:

- Under Microsoft Windows, the caption of suspended isolate processes is changed to help identify the relevant process (see 4.8 on debugging).
- The StartServer function now takes a list of rules which filter the network addresses of clients that are allowed to make use of the server (see section 6.4 on running Isolate Servers).
- A number of bug fixes related to recovery following failures, better diagnostic messages, and portability to Linux and AIX.

2 Philosophical Preamble

This section contains a discussion of issues surrounding the parallelisation of APL code. If you are looking to quickly get started with using futures and isolates, skip to section 3.

2.1 APL is a Parallel Notation

The APL language contains many language features which are potentially parallel, although most implementations do not execute them in parallel:

- Most primitive functions work on arrays, performing operations on all the elements of argument(s).
- Many classical primitive operators from APL or APL2 such as each (f^*), inner product ($f \cdot g$), outer product ($\circ \cdot g$), and some reductions ($f /$) and scans ($f \backslash$), express repeated and potentially parallel execution of one or more functions.
- New operators added to Dyalog APL more recently, such as rank ($f \circ n$), key ($f \boxplus$), and even some uses of the power operator ($f \star g$), also express potentially parallel operations.
- In Dyalog APL, the use of the dot to execute an expression within the scope of an array of objects or namespaces, expresses potentially parallel execution (`namespaces . foo`).

Starting with version 14.0, Dyalog Ltd. intends to give high priority to implementing actual parallel execution of APL language statements.

There are two main tracks to pursue:

- Automatic parallelisation, where the interpreter automatically infers that an APL language statement can be parallelised, and decides to execute code in parallel when the language processor believes this will improve the performance of an application.
- Explicit parallelisation, where the language provides new mechanisms that allow the programmer to declare that certain parts of the application could or should be executed in parallel.

2.2 Automatic Parallelisation

Given the number of parallel language constructs in the APL language, it might seem that automatic parallelisation should be easy to implement. In fact, parallel execution of a number of scalar dyadic functions on floating-point arguments was introduced in version 12.1 of Dyalog APL, and some user applications did realise noticeable speed-ups. However, automatic parallelisation of traditional APL expressions faces a number of significant challenges:

- 1) Most arguments are small: parallel execution has a set-up time to start parallel threads or processes, and a synchronization cost when the results computed by separate processes are coalesced into a single array result. In typical applications, the arguments to most primitive functions only have a very small number of elements. Any potential speed-up from parallel execution will be cancelled out by the setup cost.
- 2) Cost of memory access: although it is common for modern hardware to have multiple cores, the memory cannot fully support many cores reading or writing significant volumes of data simultaneously: a queue will form. Thus, operations which have a large number of

small operations on large quantities of data (such as adding or multiplying large arrays), cannot be parallelized efficiently on current hardware due to memory access costs.

- 3) Side-effects: when operator expressions are applied to user-defined functions, those functions can have side-effects. Many existing applications rely on the current order of execution of operator expressions and would fail or produce incorrect results if user-defined functions started executing in parallel.

Successful automatic parallelisation on existing hardware will require an APL language system which is able to coalesce multiple primitive APL operations into larger units which can be effectively executed in parallel – combined with an ability to detect that user-defined functions are free from side-effects (in other words, a “compiler”).

Dyalog Ltd. is conducting and funding research into APL compilers and does expect to increase the amount of automatic parallelisation that Dyalog APL supports. However, in the short term, explicit parallelisation through futures and isolates promise to provide more “bang for the buck” for the application developer.

2.3 Explicit Parallelisation

As mentioned in the previous section, it can be difficult for the APL engine to determine whether parallelisation is both safe and worthwhile. However, in many cases the programmer knows where the potential for significant parallel execution lies in an application, and is able to identify sections of code that are free of side-effects (or have manageable side-effects) - and represent enough work for the setup and synchronization costs to make parallel execution worthwhile.

In the past, many languages have provided language extensions aimed at allowing parallel processing¹. These have typically depended on the programmer to make use of explicit synchronization features such as semaphores, which allow threads to mutually exclude each other from interfering with each other as they modify shared data, notify each other of progress, and wait for one step to complete before a dependent task can be started. These features are notoriously difficult to use and often give rise to deadlocks and other difficult-to-detect defects in code, even when used carefully by quote trained professionals unquote.

The challenge is to come up with new language features which make it easy for APL users to introduce parallelism without reducing the readability of the code – and without making the application fragile and prone to synchronization of timing errors.

This has been the main design goal for futures and isolates.

¹ Including Dyalog APL, which provides a spawn operator (&), the `⊖TGET/⊖TPUT/⊖TSYNC` family of system functions, and the `:Ho l d ... :EndHo l d` control structure.

3 Introducing Futures and Isolates

Futures and Isolates are language extensions which are intended to provide the developer with ways to express the existence of potentially parallel sections of code in a *deterministic* fashion. Briefly:

- An *isolate* is a namespace which is semantically equivalent to a normal namespace, except that expressions which are executed inside an isolate can run in parallel (in a separate process) to the main application process². The current proposal is to create isolates using a new primitive function α (*isolate*).
- A *future* is an array of unknown rank, shape and content, which is returned as the result of any expression that is executed inside an isolate. Futures can be passed as arguments to user-defined functions or “nested” to form arrays containing several futures, without blocking (waiting for the isolate to produce a result). Arrays containing futures can be subjected to structural transformations such as reshape or partitioned enclose, without blocking. Blocking only occurs when one or more futures are passed to a function which needs to know the actual value (for example, a mathematical primitive function).

Between them, futures and isolates make it straightforward to divide application code into sections which can run in parallel: if one function requires the result of code which is still executing in parallel, the dependent function will automatically wait until the data that it needs is available.

3.1 Isolates

An isolate created using α is very similar to a namespace created using \square NS. The contents of an isolate can be referenced using the same dot notation that is used to refer to the contents of a regular namespace. It is also possible for code running inside an isolate to refer back to the parent space using the ## (Parent) symbol. However, isolates reside in separate processes and do not share a common workspace with the main process, and this does introduce some differences.

The most important differences between isolates and namespaces are the features that make isolates a tool for simple implementation of parallel execution within an application:

- Any evaluation which occurs within an isolate is handled in a separate process: If more than one isolate exists, execution can proceed independently within each isolate. If more than one processor is available, expressions inside isolates will run in parallel with the main process.
- An expression executed within an isolate immediately returns a *future*, without waiting for the expression to finish execution. A future is a placeholder for an as yet un-computed

² This is different from the threading that results from the use of the & (spawn) operator: These threads all run inside a single APL process, and only one thread is actually running at any given time.

value (see section 3.2 for more information). The future automatically turns into a real value when the expression produces a result.

- Errors and interrupts are treated differently: In particular, errors are not reported to the calling process until the value is referenced. Note that this means that, if values are never referenced, errors in the code that (would have) produced them will go un-detected.

A number of differences are due to the fact that the isolate is actually inside a separate process. Although the contents – both code and data – are “only a dot away”, there are some restrictions on making references between isolates (or they would not deserve the name).

- Because the isolate is running in a separate process, it does not share file ties or any process-related handles or resources with its parent process³.
- Inside an isolate, the special symbol `##` (aka “Parent”) refers to the root (`#`) of the parent process - and not, as one might expect, the space inside which the isolate was created (as would be the case for a normal namespace).
- Code running inside an isolate can *only* refer to spaces that are not contained within itself *via* the root of the parent workspace (referenced using `##`, as mentioned above). You cannot pass a reference to something contained within one isolate to another isolate.

A handful of differences are due to the current implementation. These may disappear or change as we collect feedback and the implementation improves and eventually becomes completely integrated with the interpreter. Examples of these differences include:

- Expressions which produce a shy result when executed in a namespace will no longer be shy when executed in an isolate, because a future was returned and the shy nature of a result cannot presently be preserved by futures. For example: `nsref . (X←42)`.

For about the differences between isolates and namespaces, see section 7.

3.2 Futures

A future is an array of unknown rank, shape and content, which is returned as the result of any expression that is executed inside an isolate. Futures have the following characteristics:

- A future can be assigned to a variable, passed as the argument to a function or operator, stranded together with other elements to form an array, or be inserted into a nested array, while remaining a future (with unknown content).
- Arrays *containing* nested (enclosed) futures can be passed as arguments to user-defined functions or operators, or subjected to primitive structural functions such as `reshape` or `compress`, so long as no function which needs to know the shape or content of the future is encountered.
- If a primitive or system function which needs to reference the contents of the future is encountered, an attempt is made to display the future in the session, or pass it to an external program such as a Microsoft.NET method or a shared library function, the

³ In fact, as we shall see later, an isolate *could* be running on a completely different machine, which might not even be running the same operating system or be able to see the same file system as the parent process.

current thread will suspend until the expression which produced the future either produces the result, or the expression fails.

In other words, future results of expressions that are currently being executed inside one or more isolates (which are running in separate processes) can be passed around an application until the actual values are needed for the next computational step, at which point the code (the thread) that needs it will automatically block until the value is available.

Together, isolates and futures are designed to make it straightforward to write applications which contain sections that can run in parallel. No semaphores or special mechanisms are required to synchronize the independent expressions: When a result that is being computed asynchronously is required, the consumer of the result will simply wait until the value becomes available.

3.3 Parallel Operator

Parallel (//) is a monadic operator which returns a derived function which will immediately return a future, and execute the operand function within an empty isolate; it provides a mechanism for parallel execution of a function without the overhead of explicitly creating an isolate using α .

The parallel operator is typically combined with each (**), in which case the operand function is invoked multiple times. For example, the elapsed time for the following expression should be slightly more than 3 seconds (without the parallel operator, it would be roughly 12 seconds):

```

□DL ||" 3 3 3 3
3.003 3.004 3.003 3.003

```

The temporary isolate is empty when the function starts running; the operand function may create global names for its own use, but they will be discarded together with the temporary isolate when execution completes.

4 Introducing the v14.0 Implementation

The ability to deal with the new type of array known as a future is a primitive feature of Dyalog APL version 14.0, built-in to all primitive functions. However, the mechanisms for the creation of isolates and the execution of expressions within them are currently modelled in APL with the help of an undocumented I-Beam.

The model can be found in the distributed workspace `ws\isolate`, in the form of a namespace called `isolate`, and a set of functions and operators that have been given names that are intended to be similar to future primitive expressions. The following functions and operators expose models of future primitive functions or operators:

Syntax	Primitive Equivalent	Alternative Name	Description
\emptyset arg	\varkappa arg	<code>isolate.New</code>	Creates a new isolate
f II	f	<code>isolate.ll</code>	Parallel operator
f Iİ	f "	<code>isolate.llEach</code>	Equivalent of Parallel Each
f o_II	$\circ.(f)$	<code>isolate.llOuter</code>	Parallel outer product
f IIö ranks	$(f)ö$ ranks	<code>isolate.llRank</code>	Parallel rank operator ⁴
f IIĐ	$(f)Đ$	<code>isolate.llKey</code>	Parallel key operator <small>(see rank footnote)</small>

In the rest of this document, most parallel examples will make use of the above shortcuts: In particular, \emptyset will be used in place of future \varkappa primitive and Iİ in place of `||"`. Note that alternative, more traditional names are available for the names exposed in the root namespace.

The `isolate` namespace also exposes a set of functions which can be used to set configuration options, query the state of the isolate framework, start and use "remote" isolate servers. These functions will be introduced in different sections of this document, and are listed in Appendix A.

4.1 Disable "Pause on Error"

Under Microsoft Windows only: The isolate framework runs a number of threads in the background and, if "Pause on Error" is enabled, then any un-trapped error caused while experimenting is likely to freeze the system until you use the "Resume all Threads" menu item.

Rather than remembering to resume threads if anything seems to be taking too long, it is probably easier to disable pausing on error using the item in the "Threads" menu.

⁴ In the current model, to avoid blocking in order to "mix" the results into the complete result matrix, IIö (rank) and IIĐ return nested arrays containing futures, which need a final "mix" to produce the same result that the primitive operator would have returned. See Appendix A for details.

4.2 Parallel Each

A trivial example, which illustrates the potential of isolates very well, is the behaviour of parallel invocations of the system function `⊞DL`, which consumes almost no system resources, but takes a fixed amount of elapsed time. In “normal” Dyalog APL:

```
T←⊞AI ⋄ ⊞⊞DL¨8ρ3 ⋄ ⊞⊞AI-T
3.008 3.004 3.005 3.004 3.007 3.009 3.007 3.007
0 7125 24054 0
```

The second line of output shows that the elapsed time for this expression is very slightly more than 24 seconds, which should be no surprise: the each operator applies the `⊞DL` function sequentially to each of the 8 elements in the right argument.

In Dyalog v14.0, we are able to run the 8 delays in parallel using the operator named `⊞I`, which represents the proposed primitive operator parallel used in conjunction with the each operator: `//¨`:

```
T←⊞AI ⋄ ⊞DL ⊞I¨ 8ρ3 ⋄ ⊞AI-T
3.002 3.002 3.002 3.004 3.004 3.003 3.004 3.003
0 3000 3051 0
```

We can see from the timings that there is a slight overhead due to the cost of creating the isolates and communicating with them – a small number of milliseconds per function invocation.

Calls to `⊞DL` parallelise exceptionally well, since they consume virtually no resources – but they are perhaps not THAT useful in practice. Let us examine what happens if we use a (very slightly) more interesting example. Consider the following function which loops a given number of times, does some fairly hard (if useless) work, and returns the number of elapsed milliseconds as its result:


```
▽ r←loop n;work
[1]   r←3⊞AI
[2]   :Repeat ⋄ n+n-1
[3]   work←≠u10000?10000
[4]   :Until n≤0
[5]   r←(3⊞AI)-r
▽
```

Unlike `⊞DL`, this function consumes a large amount of CPU time – it is in a tight CPU loop. Let us repeat an experiment similar to the one above:

```
T←⊞AI ⋄ loop¨ 4ρ1E4 ⋄ ⊞AI-T
2958 3036 3039 3002
0 12047 12040

T←⊞AI ⋄ loop ⊞I¨ 4ρ1E4 ⋄ ⊞AI-T
6561 6621 6631 6611
0 5829 6664 0
```

The following Windows Task Manager screen shot shows the profiles of these two executions: The first run taking 12 seconds at about 30% reported utilization, the second 6 seconds at 100%.



We can see that the parallel invocations come fairly close to running twice as fast, completing the four calls to `loop` in 55% of the time (6.6 vs 12.0 seconds). This is not a bad result on the machine where this was run, a fairly typical laptop of the current era – with a 1.6GHz Intel Core i5-4200. The machine is reported as having 4 logical processors, but there are only two independent cores.

4.3 Explicit or Persistent Isolates

When a function is invoked under one of the parallel operators, an empty isolate is created for each function invocation, and the isolate is destroyed when execution completes. If you need to create persistent state that needs to allow several function invocations to share the same code or data, it makes sense to create one or more isolates ahead of time, and use them to execute more than one expression.

Isolates are created using the function `∅`, which represents the future primitive function `⌘`. The right argument can be a namespace to clone, or an empty vector to create an empty namespace. For example:

```

ns←∅NS ''           A Create an empty namespace
ns.X←1 2 3 4       A Give the name X a value within the NS
is1←∅ ''          A Create an empty isolate
is1.X←5 6 7 8 9   A Give X a value there
is2←∅ ns          A Clone the namespace as an isolate
isolates←is1 is2  A References to isolates can form an array
isolates.(+/X)    A Execute (+/X) in 2 isolates in parallel
10 35
T←∅AI ∅ isolates.(∅DL 3) ∅ ∅AI-T
3.003 3.003
0 3000 3010 0

```

Since the dot notation allows us to execute any expression within an isolate, we can also easily transfer code to an isolate “on demand”. We can repeat the example with our `loop` function.

```

isolates←isolates,∅'' '' A Add another 2 isolates
isolates.∅FX c∅CR 'loop' A Fix loop function in each isolate
loop loop loop loop
T←∅AI ∅ isolates.loop 1E4 ∅ ∅AI-T
6714 6783 6801 6788
0 5890 6814 0

```

Performance note: calls to existing isolates are generally very slightly faster, since no isolate needs to be created first. The above example happened to take a little longer runtime than the earlier example using `II`, but this will have been due to “other things going on” in the background.

4.4 Blocking

An expression executed within an isolate *immediately* returns a future, while the expression continues to execute in the background. Arrays containing futures can be manipulated using structural primitives, and passed as arguments to user-defined functions. Your application will only wait (or *block*) for the result if you encounter an expression which needs to know the actual *value* of the future.

This is hard to illustrate in a printed document, so it is recommended that you perform the following experiments in an APL session:

```

delays←⊞DL Iİ 3 10 20 ⌘ Returns immediately
2↑delays                ⌘ Returns after about 10 seconds
3.002 10.006
delays                  ⌘ Waits until 20 seconds have passed
3.002 10.006 20.021

```

As can be seen in the preceding example, each item of the resulting array is a completely separate future, and blocking only occurs if you reference an item which has not yet been produced.

In section 4.6, we will learn how to track the execution and know which values are now available, in case we need to write applications which are guaranteed not to block (this might be a requirement for an application which provides services to others and needs to remain responsive at all times).

4.5 Errors in Asynchronous Expressions

If an asynchronous expression encounters an error, this is not revealed until an attempt is made to use the value. In the following slightly convoluted example, we invoke a function three times: the first function call will succeed, the second will fail with a `DOMAIN ERROR` after approximately ten seconds, and the third with a `LENGTH ERROR` after about 30 seconds:

```

reciprocals←5 10 30 {z←⊞DL α ⋄ 1 2÷ω} Iİ 2 0 (3 4 5)
⊞DL 15                ⌘ Wait for more than ten seconds

reciprocals
FUTURE ERROR: 11: DOMAIN ERROR: {z←⊞DL α ⋄ ÷ω}
reciprocals
^

reciprocals[1] ⌘ This is fine
0.5 1

reciprocals[2] ⌘ This failed after 10 seconds
FUTURE ERROR: 11: DOMAIN ERROR: {z←⊞DL α ⋄ ÷ω}
reciprocals[2]
^

```

```
reciprocals[3] # Fails after 30 seconds
FUTURE ERROR: 5: LENGTH ERROR: {z←[]DL α ◊ 1 2÷ω}
reciprocals[3]
^
;[]DMX.(EN EM Message)
```

86
FUTURE ERROR
5: LENGTH ERROR: {z←[]DL α ◊ 1 2÷ω}

Important points:

- 1) Although the first error occurred after 10 seconds, while the main process was performing a (`[]DL 15`), no error was reported at that time. The error is only when a reference is made to a failed future. Note that this means that if you never reference the result, the error will never be reported⁵.
- 2) If we are blocking on several futures, an error is reported as soon as the *first* one of them fails.
- 3) Errors are directly associated with the corresponding future: Note that the errors reported for item [2] and [3] are different (`DOMAIN ERROR` and `LENGTH ERROR`), and that item [1] can be extracted despite the fact that the array contains failed futures.
- 4) Like any item of an array, a future can be replaced – see below.
- 5) All failures in asynchronous expressions signal a `FUTURE ERROR`, event number 86. Information about the underlying error is provided in `[]DMX.Message`.

Beware: The exact format of the error messages produced by failures of asynchronous expressions is quite likely to change as the implementation matures. It is recommended that you write applications in a way that allows you to easily adapt to changes in error reporting.

As mentioned above, a future can be overwritten at any time - it is just an item of an array. Thus, we can “repair” an array containing failed futures by replacing the failed future with any value we like:

```
reciprocals[2 3]←-'[failed]' # Overwrite the failed future
reciprocals
0.5 1 [failed] [failed]
```

If you overwrite a “live” future (where the isolate is still at work), the result of the asynchronous expression will simply be discarded when it eventually completes.

⁵ If a tree falls in the forest, and nobody is there, does it make a sound? In this case, the answer is “no”.

Beware: Even if you overwrite a “live” future, the isolate will keep running until it has finished the calculation that was launched. The isolate remains in a busy state, and will not be able to process another expression until the current one has finished executing. There is currently no way to stop an isolate from continuing execution of an asynchronous expression.

4.6 Tracking the Status of Asynchronous Expressions

It can be useful to know whether an array contains un-computed futures, in order to avoid blocking. In particular, any code which needs to remain responsive in order to provide services to others should try to avoid unexpected blocking. The isolate namespace provides four functions to help you monitor the status of futures without the risk of blocking. The right argument to each function must be the `name` of an array:

Function	Return Value
<code>Values</code>	An array of future values with the same size as the named array, with unfulfilled futures replaced with the value given as the left argument.
<code>Available</code>	A Boolean array with 1 marking values which are available.
<code>Failed</code>	A Boolean array with 1 marking futures which have encountered errors.
<code>Running</code>	1s identify futures where the isolate is still running.

For example:

```
z←99, □DL Iİ 3 6 9 -1
□DL 7 A Wait until an “interesting” point in time
isolate.((-1 Values'z'),(Available'z'),(Running'z'),-Failed'z')
99      1 0 0
3.009 1 0 0
6.015 1 0 0
-1      0 1 0
-1      0 0 1
```

Explanations:

Values: the first item (99) was never a future so the value is immediately available. 3.009 and 6.015 are the first two □DL results. -1 is returned for “unfulfilled” futures.

Available: the first three values are available, the last two are not.

Running: the fourth item is a future which is still in process after 7 seconds (the □DL 9).

Failed: the last item of `z` is a failed future, because -1 is not a valid argument to □DL.

4.7 Interrupting Asynchronous Expressions

If you issue an interrupt while your application is blocked on a future, you will be asked to confirm whether you would like to continue waiting. In the first example below, we respond `Y`, resumed waiting and received the result. At the second prompt we answered `N`: in this case the future is treated in the same way as if the calculation had encountered an error.

```
□DL Iİ 3 6
ISOLATE: Interrupt - continue waiting (Y/N)? Y
3.003 6.009
```



```

      ⎕←z←⎕DL IÏ 3 6
ISOLATE: Interrupt - continue waiting (Y/N)? N
FUTURE ERROR: USER INTERRUPT 1002
      ⎕←z←⎕DL IÏ 3 6
      ^

      isolate.Values 'z'
3.006 [Null]

```

The final result shows us that the first item of the result was produced before we interrupted the wait.

Note that interrupting a wait for a future does NOT interrupt the code running within the isolate. The computation will continue and the result, when it is available, will be discarded. Since calls to isolates are processes sequentially, the next calculation will be held in a queue until the previous function call is completed.

4.8 Debugging Asynchronous Expressions

As we have seen, the default behaviour is to trap all errors occurring within expressions that are executed in an isolate, and return the error to the calling environment:

```

      ⎕FX 'R←A CALC B' 'R←A÷2×B'
      1 CALC 2
0.25
      1 CALC IÏ 2 3 0
FUTURE ERROR: 11: DOMAIN ERROR: CALC[1] R←A÷2×B
      1 CALC IÏ 2 3 0
      ^

```

During development of code which is to run in isolates, trapping of errors can make debugging rather tedious. Fortunately, we can set the `onerror` option to `debug`:

```

      isolate.Config 'onerror' 'debug' ⌘ Returns previous value
signal

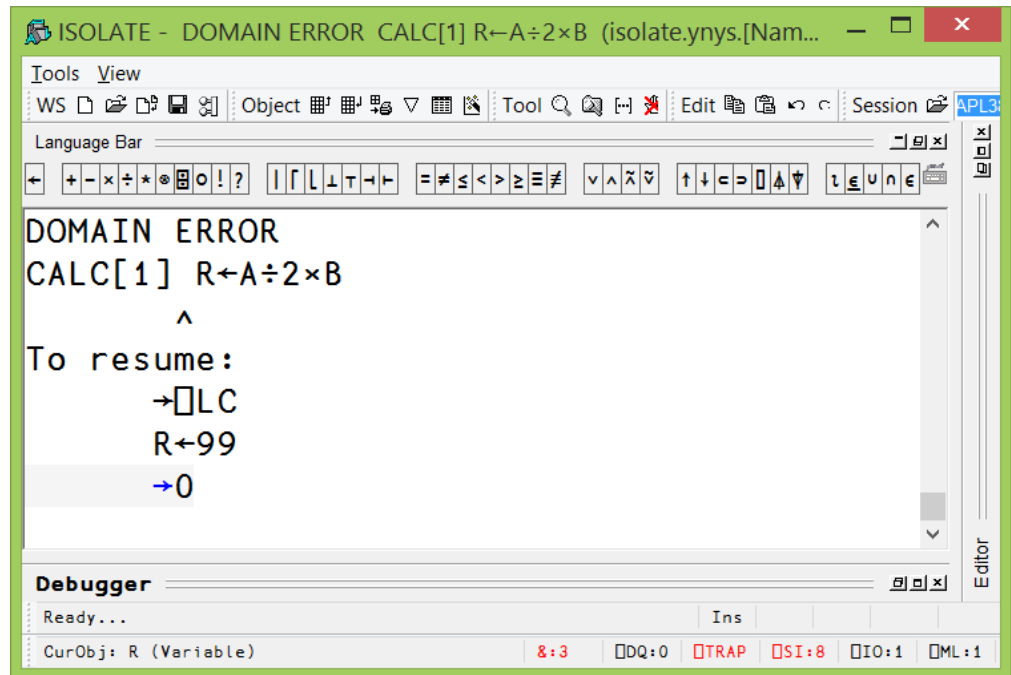
      isolate.Reset 0 ⌘ All processes must be restarted6
Reset: 2 isolates, 6 processes

      1 CALC IÏ 2 3 0

```

... at this point the APL session will hang indefinitely, due to one of the isolate processes being suspended with an error. Selecting `debug` overrides the runtime option, forcing the use of the full development environment to host isolate processes. The session caption for failed sessions should have been changed to make it easy to identify which session to select. This is a completely normal Dyalog APL development session: if you wish to resume normal operations after you have corrected the error, you must take care not to cut the stack back past your function (or if you do, try to carefully trace back into it). Ideally, you should fix your problem and continue execution of your function in such a way that it returns a result. If everything comes unstuck, return to the main session and call (`isolate.Reset 0`) again; this will terminate all the processes.

⁶ A reset is necessary after most configuration changes, because isolate processes need to be restarted in order to pick up the new configuration.



We can return to the client session and see that the result has arrived:

```
0.25 0.1666666667 99
```

Of course, you should ensure that any code changes you make to your own application code are saved in your source code repository before you return control to the isolate framework.

5 Calling Back from Isolates

In all the examples we have seen so far, we have issued calls from the main process into isolates – either to persistent isolates created using `∅`, or to temporary isolates created using one of the operators like `Iİ`.

5.1 Calling Back to the Parent

Sometimes, there are resources which it is impractical to copy to into each isolate, either because they change frequently and all isolates need to share the values in real time, or because the cost of replicating the resource make it impractical. For example, if the main process has an open database connection, there might be good reasons for isolates to call back and share that connection (or current transaction), rather than making a new database connection from each isolate.

For code running in an isolate, the special `##` (parent) is a reference to the root namespace of the initiating process, and any reference to the contents of `##` represent a call back to the main process. In the APL model, call-backs are not enabled by default, because there is significant overhead in establishing the additional communications channel that is required.

The `listen` property is used to enable callbacks:

```

    isolate.Config 'listen' 1 A Enable callbacks
0
    isolate.Reset 0 A Restart all services after config change
0 isolates, 4 processes reset

```

Imagine that we have a shared counter that all isolates need to increment from time to time:

```

    COUNTER←1
    {##.COUNTER←##.COUNTER+ω} Iİ 1 1 1 1
2 2 2 2
    COUNTER
2

```

Our intention was that each parallel function would have incremented the counter once, but something went wrong. The reason is that each expression makes two references to `##.COUNTER` – first to read it and then to set it, after adding 1. In the above example, all the four reads happened in parallel, and as a result each isolate got hold of the value 1 before any of the increments had taken place. Thus, all four added 1 and set `COUNTER` to 2. We need a way to ensure that each function manages to read and update `COUNTER` without interference from the other functions which are running in parallel.

It is not possible to use tokens (`∅TPUT` / `∅TGET`), or `:Hold` controls structures to achieve this, because the isolates are separate processes and do not share token pools or any other state with each other.

The simplest way to implement this kind of *transaction* is to use a server-side function. All call-backs to the main process are guaranteed to run sequentially. The same is true if several threads in the main process make calls to the same isolate: if there are simultaneous calls they will be queued and executed one by one, to avoid race conditions.

Let us define a function in the root namespace to modify COUNTER:

```

    ▽ R←INCR X
[1]   COUNTER←COUNTER+X
[2]   R←COUNTER
    ▽

```

Now, we can make atomic updates:

```

    COUNTER←1
    {##.INCR ω} Iİ 1 1 1 1
2 5 3 4
    COUNTER
5

```

Note that the result illustrates that the four call-backs were not run in order from left to right: the order of execution of isolates is not guaranteed – they run completely independently. The important thing is that COUNTER was incremented four times – as expected.

5.2 Call-backs versus Interrupts

Beware: Dyalog APL does not currently provide a way to select which thread is interrupted (this will probably be added to version 14.1). If you have enabled call-backs using the “listen” option (see section 5.1), and you interrupt execution, there is a risk that you will interrupt a call-back, or terminate the listening thread while it is waiting for the next call-back.

If you unintentionally interrupt the listener, you will see warnings displayed in the session, similar to the ones coloured blue in the following. The first one is a warning from the listener that it is shutting down, and the second message tells us that the listener is being restarted:

```

    is.⊞DL 6
Soft interrupt received, RPC Server shutting down
6.011

    DIVISOR←10
    {ω÷##.DIVISOR} Iİ 1 2 3
ISOLATE: Callback server restarted
0.1 0.2 0.3

```

We can see from the above that the interrupt failed to stop the first function call, but instead it caused the listener to shut down.

5.3 Calling From One Isolate to Another

The full primitive implementation may allow calls to be chained, so that one isolate can call another by going via the parent process, as in:

```
is1.(##.is2.(⊞DL 1))
```

However, this is *not* currently supported, and it is not clear that the APL-coded model will ever be able to support it. A primitive implementation may eventually also allow an isolate to create child isolates and make calls to them, but this is also unsupported in the model.

Note that, even when calls between isolates become enabled, it will probably not be possible for one isolate to contain a reference to another – references cannot cross isolate boundaries.

6 Isolate Processes

By default, the framework will create one process for each processor that is available on your machine. As isolates are created, they will initially be allocated to free processes. If no free processes are available, each new isolate will be allocated to the process which is managing the smallest number of isolates. If several processes have the same number of isolates and some of them contain “busy” isolates, the least busy process will be selected to host the new isolate.

6.1 `isolate.State`

The function `isolate.State` can be used to monitor isolate allocation and activity:

```
isolates←⊖''6ρ<' '
isolate.State ''
Host      Port  Isolates  Busy
----      -
localhost 7052      2         0
           7053      2         0
           7054      1         0
           7055      1         0
```

We can see from the above that we have 4 processes available, each listening on a separate port, and that the two first ones have been allocated two isolates each. The “Busy” column shows how many isolates are currently running code:

```
z←⊖DL Iİ 16 ⋄ ⊖DL 3
isolate.State ''
Host      Port  Isolates  Busy
----      -
localhost 7052      2         0
           7053      3         1
           7054      2         1
           7055      2         1
```

We created another six isolates for a total of 12, so each processor had 3 each – but the temporary isolates are disappearing as each delay completes. After 3 seconds, 9 isolates remain, as documented above.

6.2 Limiting the number of Isolates per Process

By default, several isolates may share a single APL process. If you are running code with no side-effects, the only problem will be that two or more isolates share a single operating system thread and the workspace that is available within it – which may limit performance.

However, if your isolates use component files, shared libraries or other “per process” system resources, these will now be shared – not with the main process, but between “random” isolates that happen to share the same process. This may cause problems unless your code is written to deal with that situation.

If your code needs to run in a single process, you should set the `isolates` configuration option to 1. Obviously, imposing a limit creates a risk that you will run out of processes:

```

isolate.Config 'isolates' 1
99
isolates←0"4p<" A OK to create four isolates
fifth←0 ''
ISOLATE ERROR: All processes are in use
fifth←0 ''
^

```

See the next section for information on how to increase the number of processes.

6.3 Changing the number of Processes

By default, the system detects the number of processes available, and sets the `processors` option to this value. The `processes` option, which controls the number of processes to start per processor, defaults to 1. The default values are expected to give a reasonable utilisation of multi-core hardware for fairly CPU-intensive isolates. However, depending on whether your application is memory- or CPU-bound, or limited by file or network I/O, the “best” choice might be very different. It might make sense to run only one or two processes – or dozens – depending on what your isolates are going to be doing.

You can set these two options to any integer value – and then you must call `isolate.Reset` in order to get the new settings to take effect. Note that in the current version, every one of the (`processors×processes`) APL tasks are started as soon as the isolate framework is initialised or reset. A potential future enhancement would be to allow dynamic allocation.

```

isolate.Config 'processors' 3
4
isolate.Config 'processes' 2
1
isolate.Reset 0 A Get new config to take effect
Reset: 6 isolates, 4 processes

z←⊖DL IĪ 16 ⋄ ⊖DL 2
isolate.State ''

```

Host	Port	Isolates	Busy
localhost	7052	0	0
	7053	0	0
	7054	0	0
	7055	1	1
	7056	1	1
	7057	1	1

6.4 Running Isolate Servers

By default, isolates are hosted in processes that are started on your local machine, according to the `processors` and `processes` options described in the previous section. If you have processing capacity available on other machines, you can load the `isolate` workspace on these machines and use the function `isolate.StartServer` to turn the current APL session into an isolate server which can be used to run slave processes.

The right argument to `StartServer` is a list of access rules, separated by commas, which control which client machines will be allowed to connect to the server. If the right argument is empty, the server will only be usable from clients on the same machine.

Each rule is of the form `'keyword=value'`. At the moment, the only keyword which is accepted is `"ip"`. The value must be a full or partial IP address. If the address is a prefix of the address of a client, the connection will be accepted. For example, `'ip=192.168.1'` would allow all machines with an address starting with 192.168.1 to connect.

The number of processes started is sensitive to the same option settings as for local use:

```
isolate.Config 'Processors' 8
4 isolate.StartServer 'ip=192.168.1,ip=192.168.6'
Machine Name: MortensYoga
IP Ports:      7052 7053 7054 7055 7056 7057 7058 7059
```

Enter the following in the client session:

```
#.isolate.AddServer 'MortensYoga' (7052-IO-t8)

Full IP address list:
IPv4   192.168.193.1   IPv6   [fe80::545f:770f:dab:368b]
       192.168.247.1   [fe80::ccbd:aa34:3cfb:6e24]
```

You can use `(isolate.Reset 0)` to shut the server down again – or simply end the session.

6.5 Using an Isolate Server

As suggested in the output of the `StartServer` function, you need to call `AddServer` to make a remote server available as a host for new isolates. The argument is a 2-element vector containing the IP address of the server, and the range of port numbers that is exposing. For example:

```
#.isolate.AddServer 'calcsrv' (7052 7053)
Host      Port  Isolates  Busy
----      -
calcsrv   7052      0         0
          7053      0         0
```

If a server that you are using is shut down, or the connection to it is lost, errors can occur without warning if you continue using isolates hosted on a server which is no longer accessible. You may see failures within the isolate infrastructure, along the lines of:

```
DL I 14
ISOLATE: Connection to calcsrv:7053 failed: 1111
ERR_CONNECT_DATA /* Could not connect to host data port */
DL I 14
^
```

The same kind of errors *may* occur if your local isolate servers are terminated for any reason, but this is much less likely than the loss of a network connection.

6.6 Combining Local and Remote Isolates

As we have seen: if you start creating isolates without calling `AddServer`, local processes are started. If you call `AddServer` after the first use of isolates, the new processes will be

added to those which are already accessible. You may also explicitly add local processes using `AddServer` with an empty vector as the right argument:

```

      #.isolate.AddServer ''
Host      Port  Isolates  Busy
----      -
calcsrv   7052      0         0
          7053      0         0
localhost 7052      0         0
          7053      0         0

```

In the above example, the `calcsrv` (which we added in the previous section) - and the local machine both have 2 processes. Both machines are using the same port numbers (on different IP addresses).

The function `RemoveServer` can be used to remove a server from your list of remote servers. Note that this function will allow you to remove processes which are hosting persistent isolates. If you subsequently attempt to use such isolates, an error will result:

```

      is←∅ '' # Created on calcsrv
      #.isolate.RemoveServer 'calcsrv'
Host      Port  Isolates  Busy
----      -
localhost 7052      0         0
          7053      0         0

      is.(2+2)
ISOLATE: No longer accessible
      is.(2+2)
      ^

```


7 Differences between Namespaces and Isolates

This chapter drills down on some of the less obvious differences between isolates and namespaces that may be encountered when converting existing code to make use of isolates.

7.1 File Ties are not shared with Isolates

Unlike namespaces, isolates reside in separate processes, so they do not share any resources that belong to the current process, such as file ties, loaded libraries, TCP sockets, etc. The following example illustrates the issue using component file ties:

```

ns←NS '' ◇ is←∅ ''  A A namespace and an isolate

'file1' []fstie 0    A Tie a file in main process
1
'file2' ns.[]fstie 0 A Tie a file "in" a namespace
2
'file1' is.[]fstie 0 A Tie first file in the isolate
1

ns.([]fnames,;[]fnums) A root and ns are in the same process
file1 1
file2 2

is.([]fnames,;[]fnums) A isolate is a separate process
file1 1

```

7.2 Namespace “Refs” and Isolates

In addition to the resources described in section 7.1, it is also not possible to share pointers such as namespace refs with an isolate. The process of passing the reference to an isolate, either through assignment or as an argument to a function, will cause the referenced space to be copied.

```

ns←NS '' ◇ ns.Value←42 A Create namespace containing variable
is←∅ ''              A Create an isolate
otherns←NS ''        A Create a 2nd namespace
otherns.ns←ns        A Within same process: create a ref to ns
is.ns←ns              A Separate processes: ns is copied

ns.Value←43          A Update value
otherns.ns.Value     A Same process: shared value
43
is.ns.Value           A Separate process: copied value
42

```

7.3 Classes versus Isolates

As mentioned in 7.2, namespace references cannot be shared, data will be serialised and copied when passed between isolates. Classes, Interfaces and Instances cannot be serialised, so they cannot even be copied:

```
is.MyClass←MyClass
ISOLATE: Transmission failure: Object oriented namespace cannot
be serialised
```

You can fix the source of a class in an isolate using `FIX`, but beware: if you execute an expression that returns a class or an instance, this will fail because that result cannot be passed back:

```
FUTURE ERROR: 11: ISOLATE ERROR: Result cannot be returned from
isolate
is.FIX [SRC MyClass
      ^
```

The solution is to ensure that you return a result which is not a class, instance or interface. For example:

```
is.{[SRC FIX ω][SRC MyClass A Count lines of source
21
```

7.4 Unexpected output of Shy Results

Futures are not currently able to preserve the shy nature of results – any shy result which is returned as a future becomes non-shy:

```
ns←NS ''           A Create an empty isolate
is←isolate.New ''  A Create an empty isolate
ns.(X←10?10)       A Shy result in ns: output suppressed
is.(X←10?10)       A With an isolate, this produces output
6 9 3 4 8 2 7 1 5 10
```

The parenthesised assignment within a namespace in the current process is shy and does not produce any output; when the same assignment is made within an isolate, a future is returned and the result of the assignment is displayed in the session.

7.5 Function calls into Isolates must return results

Due to the way that calls into isolates are implemented as remote procedure calls that return futures, any expression executed in an isolate MUST return a result, or a FUTURE ERROR will be signalled.

Note that when an error is signalled because there is no result, or because the result is a class, function or anything else that cannot be returned, the expression may well have been executed.

8 Option Settings

The `isolate` namespace has a number of configuration options which control how isolates are started and managed. Options are queried and set using the `Config` function, for example:

```
isolate.Config 'workspace' 'c:\temp\myapp\myapp.dws'
c:\temp\myapp\myapp.dws
```

Most of the options will only take effect following a call to (`isolate.Reset 0`) in order to shut down all existing servers and re-initialise the system.

An empty argument to the `Config` function gives a list of all current settings:

```
isolate.Config ''
drc                #
homeport           7051
homeportmax        7151
isolates           99
listen             0
maxws              64000
onerror            signal
processes          1
processors         4
protocol           IPv4
runtime            1
workspace          isolate
```

The current set of supported options is listed below.

Option Name	Default	Description
<code>drc</code>	<code>#</code>	Location of CONGA namespace to use
<code>homeport</code>	7051	The lowest port number that will be used
<code>homeportmax</code>	7151	The highest port number to try listening on
<code>isolates</code>	99	Number of isolates allowed per process
<code>listen</code>	0	1 to allow isolates to issue callbacks to parent process
<code>maxws</code>	'64000'	By default, uses the same setting as the current APL session
<code>onerror</code>	'signal'	Signal errors to the line waiting for results
<code>processes</code>	1	The number of processes to start per process
<code>processors</code>	4	Number of processors (under Windows, determined automatically)
<code>runtime</code>	1	Whether to run isolates using the runtime engine
<code>workspace</code>	'isolate'	Workspace to load when starting new isolates

8.1 drc

This option can be used to specify a non-default location of the Conga DRC namespace, which must be present for isolates to work. If `drc` is `#` (the default) we check to see if `#.DRC` exists and copy DRC from the CONGA workspace (which must be on the workspace path or in the same folder as the APL executable) if it not. Otherwise `drc` must be a reference to an existing DRC namespace.

8.2 homeport and homeportmax

Each isolate - and if the `listen` option is set to 1 the main process as well - will open a listening TCP port in order to receive commands to be executed. These two options control the first and last port number that will be attempted used. The system will use the configured number for the main process listener, and the next $(\text{processors} \times \text{processes})$ numbers for the isolates.

If the creation of the main process listener fails, the port number will be incremented by $(1 + \text{processors} \times \text{processes})$ and the listener creation will be retried - up to the upper limit set by `homeportmax`.

8.3 isolates

This parameter makes it possible to set a cap on the number of isolates that each available process is allowed to contain (within a process, separate APL threads will be used to execute code for different isolates). By default, the limit is 99 - but if you are running isolate code which needs to have complete control of process resources such as component files or shared libraries, you can set this value to 1.

8.4 listen

When this option is set to 1, a listening thread is launched in the main APL process, and the isolate will be able to call back into the root of the main workspace by referring to the symbol `##`. The default is 0 because there cost of creating isolates is significantly increased when an additional connection is required. The cost of using isolates, once they have been created, is not affected.

8.5 maxws

If you want to start processes using a different workspace size than your main process, you can set this option before processes are launched. The option value is a string which has to follow the same rules as the `maxws` environment variable or registry entry for APL.

8.6 onerror

This option controls the handling of errors that occur in expressions that are executed in isolates. Allowed values are:

`signal`: a FUTURE ERROR will be signalled in an expression which depends on the result of the expression.

`debug`: do not trap errors in the isolate: let the isolate server suspend. This will force the use of the full development environment for isolate processes, even if the `runtime` option is set to 1.

8.7 processors

The number of processors available for use for isolate servers (default determined automatically).

8.8 processes

The number of processes that will be started per available processor (default 1).

8.9 workspace

This specifies the name of a workspace to be used at startup by the processes which will host isolates. By default, it is set to “`isolate`”, and the system will use the distributed workspace by that name. If you wish to use a workspace which contains your own application infrastructure, you need to make sure that it contains the `isolate` namespace from the distributed workspace. You should also include the `DRC` namespace from the Conga workspace, or place the Conga workspace in a location where the expression (`'DRC' □CY 'ws\conga'`) will succeed.

The latent expression (`□LX`) in the workspace must start the listener by invoking the expression (`isolate.nys.isoStart θ`). Note that, if you use the runtime interpreter, you should either use the full path to the workspace or place the workspace in the same folder as the Dyalog executable.

8.10 runtime

Controls whether the runtime interpreter or the full development system is used to run isolate server processes. The default is `1` (use the runtime), which gives faster startup of isolates and no visible “clutter” caused by additional sessions. Note that the development system will be used if the `onerror` option is set to `debug`, regardless of the `runtime` setting.

9 Further Reading

9.1 Tools and Samples on GitHub

The Parallel folder in the Dyalog GitHub repository at <http://github.com/Dyalog/Samples/> contains tools that can be used to extend the functionality provided by the isolate namespace, with code examples showing how to use the tools.

At the time that this document was written (June 17th, 2014), the samples included:

PEACH – an enhanced implementation of Parallel Each which uses a fixed set of isolates and keeps them busy (rather than creating a new empty isolate for each item of the arguments), and displays an optional progress bar for the computation.

IIPageStats – an example which computes letter frequencies on a collection of newspaper web pages in parallel (using PEACH).

See the README.md files in the repository for more information.

9.2 Topics yet to come

Dyalog expects to continue to work on futures and isolates following the commercial release of version 14.0, and issue updates to both software and documentation over the summer and autumn of 2014 – as well as updating the tools and samples mentioned in the previous section.

Known issues that will be addressed in the weeks and months to come include:

- An example will be produced to show how to run isolate servers as Windows Services.
- An example will be produced to show how to use your own application workspace as the isolate server base.

Appendix A: Function and Operator Reference

A.1 Models of Future Primitives

Each syntax block shows the model syntax on the left. On the right, in the first row the proposed primitive symbol and in the 2nd row the alternative name exposed by the isolate namespace:

isolate $\leftarrow\emptyset$ init	$r \leftarrow \omega$
	isolate.New ω

Creates an isolate. The right argument can take one of the following forms:

Right argument	Meaning
' '	Empty vector: empty isolate
'wsname'	A simple, non-empty character vector is interpreted as a workspace name, the workspace is copied to create the isolate
'fn' 'var'	A vector of names is used to create an isolate containing named objects from current workspace
Namespace reference	Creates clone of argument space

The result is a reference to the isolate which has been created.

f II	f
	isolate.ll

Model of (f ||): A function derived from f, which will create an isolate empty except for a copy of the operand function, start execution of the function and immediately returns a future.

f IIX	f ''
	isolate.llEach

Model of (f ||''): A function derived from f, which will create one isolate containing a copy of f for each required invocation (using the same rules as for f'') of f, start all the function calls, and return an array of futures of the same shape as would be returned by f''.

f o_II	o.(f)
	isolate.llOuter

Model of o.(f ||): A function derived from f, which will create one isolate containing a copy of f for each required invocation (using the same rules as for o.f) of f, start all the function calls, and return an array of futures of the same shape as would be returned by o.f.

f IIö ranks	f ö ranks isolate.llRank
--------------------	---

Model of f ||ö: A function derived from f, which will create one isolate containing a copy of f for each required invocation (using the same rules as for fö ranks) of f, start all the function calls, and return an array of futures of the same shape as the outer shape of the rank invocation.

Note that, in the current model, the result is missing the final mix that would be performed by the rank operator

```

      p←(,ö 2)2 3 4p12
1 2 3 4 5 6 7 8 9 10 11 12
1 2 3 4 5 6 7 8 9 10 11 12
2 12

      p←(, IIö 2)2 3 4p12
1 2 3 4 5 6 7 8 9 10 11 12 1 2 3 4 5 6 7 8 9 10 11 12
2
    
```

The reason for the missing mix is so that an array of futures can be returned. If the mix was included, the derived function would be forced to block in order to know the shapes of the individual results.

f IIð	f ð isolate.llKey
--------------	--------------------------------------

Model of f ||ð: A function derived from f, which will create one isolate containing a copy of f for each required invocation (using the same rules as for fð) of f, start all the function calls, and return an array of futures of the same shape as the outer shape of the key invocation.


As with the model of rank, in the current model, the result is missing the final mix, in order to avoid blocking:

```

      {α ω}ð 1 1 1 2 2 3 3 3
1  1 2 3
2  4 5
3  6 7 8

      ]disp {α ω} IIð 1 1 1 2 2 3 3 3

```



A.2 Isolate Server Management

r←AddServer address ports

Prepare to use an isolate server which must already be running on a machine with the specified IP address and has started isolate processes that are listening on the specified port numbers. The result of State after the isolate processes have been added. For example:


```

# .isolate.AddServer 'calcsrv' (7052 7053)
Host      Port  Isolates  Busy
----      -
calcsrv   7052      0        0
          7053      0        0

```

An empty right argument makes local processes available:

```

# .isolate.AddServer ''
Host      Port  Isolates  Busy
----      -
calcsrv   7052      0        0
          7053      0        0
localhost 7052      0        0
          7053      0        0

```

r←RemoveServer address

Removes a server which was added using `AddServer`. The result is the state after the isolate processes have been added. For example:

```

# .isolate.RemoveServer 'calcsrv'
Host      Port  Isolates  Busy
----      -
localhost 7052      0        0
          7053      0        0

```

An empty right argument removes local processes:

```

# .isolate.RemoveServer ''
[no servers defined]

```

r←StartServer access

Starts an isolate server on the current machine. The right argument is a list of access rules, separated by commas, which control which client machines will be allowed to connect to the server. If the right argument is empty, the server will only be usable from clients on the same machine.

Each rule is of the form `'keyword=value'`. At the moment, the only keyword which is accepted is `"ip"`. The value must be a full or partial IP address. If the address is a prefix of the address of a client, the connection will be accepted. For example, `'ip=192.168.1'` would allow all machines with an address starting with 192.168.1 to connect.

It is anticipated that future versions of `StartServer` will support filtering on client SSL certificates IDs, and probably also the use of an external rule file.

The function produces output containing information that should be useful when providing arguments to the `AddServer` function:

```

isolate.StartServer 'ip=192.168.1,ip=192.168.6'
Machine Name:  calcsrv
IP Ports:      7052 7053

```

Enter the following in the client session:

```

#.isolate.AddServer 'calcsrv' (7052-IO-t2)

```

```

Full IP address list:
IPv4    192.168.193.1    IPv6    [fe80::545f:770f:dab:368b]
        192.168.247.1              [fe80::ccbd:aa34:3cfb:6e24]

```

A.3 Configuration and Process State

r←Config property [value]

Report or set configuration options (see section 8 for details). With an empty right argument, all current option settings are returned:

```

isolate.Config ''
drc                #
homeport           7051
homeportmax        7151
isolates           99
listen             0
maxws              64000
onerror            signal
processes          1
processors         4
protocol           IPv4
runtime            1
workspace          isolate

```

With a right argument consisting of an option name, the current setting is returned:

```

isolate.Config 'processors'
4

```

To set an option, follow the name with the new value. The previous setting is returned:

```

isolate.Config 'processors' 8
4

```

r←Reset 0

Resets the isolate mechanism, shutting down any isolate processes. The right argument currently has to be zero, additional options may be added in the future. The result is a brief report on the activity of the function:

```

isolate.Reset 0
Reset: 0 isolates, 4 processes

```

r←State ''

Produces a report on the current state of isolate processes. The right argument must currently be an empty character vector:

```

    isolate.State ''
[not initialised]

    is←0'' # Create an isolate: Initialises processes

    isolate.State ''
Host      Port  Isolates  Busy
----      -
localhost 7052      1        0
          7053      0        0
          7054      0        0
          7055      0        0
          7056      0        0
          7057      0        0
          7058      0        0
          7059      0        0

```

Following the headers, there is one row of output for each isolate process. The four columns of output contain:

1. The host name or IP address (argument to `AddServer`) on which the process is running. For readability, empty if the name is the same as the preceding row.
2. The port number that the process is listening on.
3. The number of isolates currently hosted by this process.
4. The number of isolates that are currently executing an expression.

A.4 State of Futures

The following functions all take the name of an array containing futures as the right argument, and return information regarding the state of each item of the array:

r←Available name

Returns a Boolean array with the same shape as the variable named in the right argument:

1: The corresponding item is not a future (either it was never a future, or the computation has succeeded).

0: The computation of the corresponding item is in progress, or has failed.

r←Failed name

Returns a Boolean array with the same shape as the variable named in the right argument:

1: The computation of the corresponding item has failed.

0: The corresponding item is not a future, is still being computed, or has completed successfully.

r←Running name

Returns a Boolean array with the same shape as the variable named in the right argument:

1: The computation of the corresponding item is in progress.

0: The corresponding item is not a future, has failed, or has completed successfully.

r←[null] Value name

Returns and with the same shape as the variable named in the right argument. If the computation is in process or has failed, the corresponding item of the result contains a `□NULL`, or the optional left argument if it was provided. Otherwise, the item contains the corresponding value from the named variable.

Appendix B: Troubleshooting

At the time that this document was written, the actual usage of isolates was essentially limited to synthetic experiments designed to test the code, or used during the writing of the document itself. Nonetheless, some problems were encountered; this section is a collection of explanations and tips to help you move forward when things stop working.

If you experience that the system frequently hangs while using futures and isolates, try the *Threads|Resume all Threads* menu item. Dyalog recommends that you do not have “*Threads|Pause on Error*” enabled when working with the experimental parallel features.

B.1 “Normal” Errors

The following error messages may appear during normal operations:

ISOLATE: Callback server restarted

The system detected that listening was enabled, but no listener process was running.

ISOLATE: Interrupt – continue waiting (y/n)?

An interrupt was issued while the system was waiting for isolates to produce results. You can choose whether to issue an error for missing values, or resume waiting.

ISOLATE: Transmission failure

An attempt was made to transmit an unserialisable value (such as a class or an instance) to an isolate.

FUTURE ERROR: 11: ISOLATE ERROR: Result cannot be returned from isolate

An attempt was made to return an unserialisable value from an isolate.

B.2 Initialisation and Communication Failures

ISOLATE: Connection to address:port failed: 1111 ERR_CONNECT_DATA /* Could not connect to host data port */

ISOLATE: Unable to connect to isolate process: PROC10584797 192.168.17.105 7052

ISOLATE: Unable to connect to started processes (attempt n of m)

FUTURE ERROR: 86: COMMUNICATIONS FAILURE 1119 ERR_CLOSED : /* Socket Closed while receiving */

These errors are most likely when using remote isolate servers: they suggest that the ip addresses or port numbers are not valid; the network has failed or the remote server has been shut down or crashed after it was added. The errors could also indicate that the started isolates are crashing immediately. A common cause of this is the inability to locate the CONGA workspace.

Suggestions:

- 1) Check whether the processes are running or the network has failed
- 2) Try switching the “runtime” property off to see whether any error messages are visible

Unable to create listener on port 7051

The system was unable to create a listener to receive callback requests. This probably means there is a defunct APL process which still has the port open. Use the “netstat” command to identify and kill it (a future enhancement to the Reset function will hopefully be able to do this automatically).

B.3 Internal Errors

The following errors indicate serious internal logic errors, and should not occur during normal operations. Please report the circumstances, ideally with instructions which reproduce the error, to the Dyalog Helpdesk:

ISOLATE: Connection OK but incorrect handshake response received

ISOLATE: New process did not respond to handshake

ISOLATE: Unknown command

ISOLATE: Internal error