

# Components for Integral Evaluation in Quantum Chemistry

JOSEPH P. KENNY,<sup>1</sup> CURTIS L. JANSSEN,<sup>1</sup> EDWARD F. VALEEV,<sup>2</sup> THERESA L. WINDUS<sup>3</sup>

<sup>1</sup>Scalable Computing Research and Design, Sandia National Laboratories, MS 9158, P.O. Box 969, Livermore, California 94551-0969

<sup>2</sup>Department of Chemistry, Virginia Tech, 403 Davidson Hall, Blacksburg, Virginia 24061-0001

<sup>3</sup>Department of Chemistry, Iowa State University, 1605 Gilman Hall, Ames, Iowa 50011

Received 9 April 2007; Revised 2 July 2007; Accepted 6 July 2007

DOI 10.1002/jcc.20815

Published online 23 August 2007 in Wiley InterScience (www.interscience.wiley.com).

**Abstract:** Sharing low-level functionality between software packages enables more rapid development of new capabilities and reduces the duplication of work among development groups. Using the component approach advocated by the Common Component Architecture Forum, we have designed a flexible interface for sharing integrals between quantum chemistry codes. Implementation of these interfaces has been undertaken within the Massively Parallel Quantum Chemistry package, exposing both the IntV3 and Cints/Libint integrals packages to component applications. Benchmark timings for Hartree-Fock calculations demonstrate that the overhead due to the added interface code varies significantly, from less than 1% for small molecules with large basis sets to nearly 10% for larger molecules with smaller basis sets. Correlated calculations and density functional approaches encounter less severe performance overheads of less than 5%. While these overheads are acceptable, additional performance losses occur when arbitrary implementation details, such as integral ordering within buffers, must be handled. Integral reordering is observed to add an additional overhead as large as 12%; hence, a common standard for such implementation details is desired for optimal performance.

© 2007 Wiley Periodicals, Inc. J Comput Chem 29: 562–577, 2008

**Key words:** electronic structure; integral; component; software development

Historically, efforts to combine functionality from quantum chemistry software packages have been limited in scope and consisted of package-specific, one-to-one solutions. Such one-to-one code integration suffers from poor scaling of programming effort; the effort to share features from  $n$  packages grows as  $O(n^2)$ . Consequently, the integration of codes from various packages is frequently considered tedious and not worthwhile, regardless of any advances in methods and capabilities, which might become available. Considering the large number of noncommercial packages existing within the quantum chemistry community, each with unique capabilities and deficiencies, this situation is unfortunate.

The future high performance computers on which quantum chemistry packages will run may depend on advanced processor architectures. Regardless of whether such systems contain coprocessors, which augment conventional general purpose processors or contain entirely unfamiliar processor architectures,<sup>1,2</sup> wide adoption of such technology will be impractical without a community code base. It is clear that a scalable approach to creating interoperable software must be adopted both to enable rapid development of advanced quantum chemical methods and to ensure such methods run reliably on the latest hardware.

Component-based software approaches break up complex tasks into loosely coupled subproblems, encouraging the definition of

standardized interfaces and enabling collaboration between research groups. The Common Component Architecture (CCA) Forum is a consortium dedicated to the development and adoption of a component architecture for scientific software.<sup>3–5</sup> Work within the CCA includes the development of standards and middleware,<sup>6–12</sup> as well as component toolkits within numerous scientific domains. Development groups adopting component technology for scientific codes span such diverse domains as optimization and linear algebra, combustion and quantum chemistry, accelerator design, fire and explosives modeling, and climate simulation.<sup>13–17</sup>

Previously, several authors participated in work towards a quantum chemistry component toolkit, using development practices

**Correspondence to:** J. P. Kenny; e-mail: jpkenny@sandia.gov

Contract/grant sponsor: U.S. Department of Energy, Scientific Discovery through Advanced Computing (SciDAC) initiative; contract: DE-AC05-00OR22725 (ORNL); contract: DE-AC06-76RLO1830 (PNNL); contract: DE-AC04-94AL85000 (SNL)

Contract/grant sponsor: Oak Ridge National Laboratory, Laboratory Directed Research and Development Program (EFV)

Contract/grant sponsor: start-up funding from the Department of Chemistry at Virginia Tech (EFV)

advocated by the CCA Forum.<sup>16–18</sup> Using CCA components to manage multiple levels of parallelism, significant improvements in machine utilization were demonstrated.<sup>18</sup> Through the development of an application for molecular structure optimization, it was demonstrated that domain scientists adopting CCA approaches are able to integrate software packages from multiple scientific disciplines.<sup>16</sup> Molecular structure optimization required the functionality of each major mathematics and chemistry package to be encapsulated in a set of components and classes. There was, however, no attempt to implement low-level interfaces between chemistry packages. Thus, while *interchangeability* was shown, substantial progress towards *interoperability* was not demonstrated.

Here we make a first step toward deeper component-based integration of quantum chemistry packages. In this work we develop a set of standard interfaces and data structures for evaluation of molecular integrals, and we demonstrate a component implementation of this design. Molecular integral evaluation is a natural low-level capability to share through components because it is the fundamental subproblem of all traditional quantum chemistry computations. As quantum chemical studies grow in sophistication, advanced capabilities, such as explicit electron correlation<sup>19–21</sup> and the inclusion of relativistic effects,<sup>22–24</sup> introduce the need for new types of integrals. The integral facilities available within each individual quantum chemistry package often lack one or more of these advanced features, limiting the range of methods, which can be implemented and made available to users of the package. Because writing efficient code for computing a new type of molecular integral requires significant development effort, it is natural to share the integral facilities as components. The obvious benefit of sharing integral capabilities between various packages is the ability to implement new theoretical methods very rapidly. For instance, in an early application of the work described here, combining integral capabilities from multiple packages allowed the latest explicitly correlated methods to take into account scalar relativistic effects.<sup>25</sup> While the scientific benefits of sharing low-level capabilities such as molecular integrals capabilities are strong, the potential performance impact of the component interface and data layout can grow beyond acceptable limits. Thus, this work serves as an important first test of the performance of CCA approaches when deeply integrated within performance-critical sections of code. Although we found the performance overhead of CCA approaches for integral evaluation to be acceptable, optimal performance does require common standards for data formats.

## The Common Component Architecture

Some of the authors have previously described in detail<sup>16,17</sup> the CCA model in the context of quantum chemistry applications. We will only briefly summarize it here.

Components<sup>26</sup> are similar to objects in that they implement some functionality and provide an interface for using it. A programmer composes applications out of objects by writing computer code, which instantiates objects and combines their functions. Although such applications can vary some of their functionality at runtime, via polymorphism, the full functionality of the application is typically determined once and for all at compile-time, i.e., statically.

In contrast, components conform to a particular environment specification, which allows composition of applications at runtime, i.e., dynamically. The runtime environment, which makes such composition possible is the component framework. The framework provides a simple (scripting or graphical) interface, which allows each end user to compose highly customized software from plug-and-play components. The CCA specification<sup>3–5</sup> has been developed expressly to meet the requirements of high-performance scientific codes, stressing high performance and respecting parallel execution.

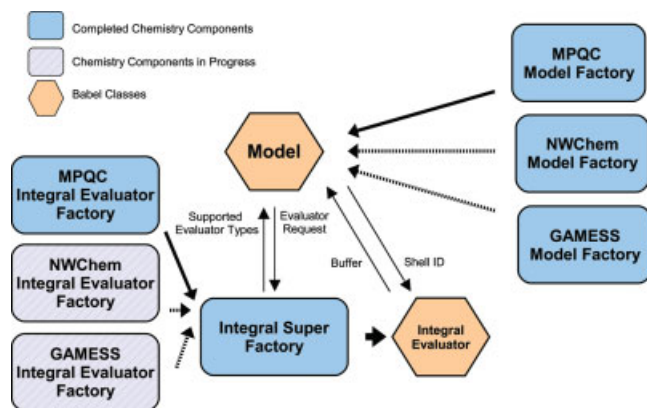
The Babel tool<sup>11,12</sup> is used within the CCA community to compose applications using components written in different languages. While not required by the CCA specification, all components described herein use Babel. Babel is a code generator, which provides implementation stubs and glue code in a collection of languages, based on interface definitions provided using the Scientific Interface Definition Language (SIDL). Babel/SIDL implements a set of fundamental data types, including complex and array types, and an object-oriented programming model for Fortran 77, Fortran 90, C, C++, Python, and Java. A SIDL interface declares methods based on these data types, and a SIDL class implements one or more interfaces. SIDL classes may be implemented in any supported language and may be used by applications written in any other supported language.

A CCA component in Babel is a SIDL class that implements the `Component` interface defined by the CCA specification and one or more programmer-defined interfaces. The `Component` interface simply defines the `setServices()` method, which is used by the framework to provide a `Services` object to the component during its instantiation. The `Services` object handles the interaction between the framework and the component. Its most important role is to inform the framework of interface implementations that the component either provides or uses. The component exposes a *provides port*, which specifies an abstract interface and provides an implementation of the interface, to the framework using the `addProvidesPort()` services method. A component can also request a *uses port* (an interface implementation) from the framework using the `registerUsesPort()` services method. Following the user's directives, the framework builds a component application by instantiating components and connecting uses ports with provides ports. Ideally, well-written software can be packaged into a component by adding a thin wrapper layer, which handles interaction with the framework. Just like the object model, the component model tends to break down as deeper levels of integration expose implementation details at interfaces and lead to the propagation of implicit dependencies throughout the application.

## Integral Component Interfaces

Here we overview the component interfaces, which we have developed. A detailed description of the interfaces is found in Appendix A.

Figure 1 illustrates the component integral evaluation architecture. The key abstractions represented by the component interfaces are the integral evaluator factory (e.g. `IntegralEvaluatorFactoryInterface` implemented by an integral super factory) and integral evaluator (e.g.



**Figure 1.** A schematic representation of the component integral evaluator architecture.

`IntegralEvaluatorInterface` implemented by each integral evaluator). Integral evaluators compute molecular integrals. Currently, we specify integral evaluator interfaces for computing integrals involving one, two, three, and four Gaussian centers (`IntegralEvaluator1Interface`, etc.). Of course, integral evaluator interfaces can be easily extended to allow more centers should such integrals prove necessary for explicitly correlated methods or other advanced techniques. Each integral evaluator provides a `compute()` method which, given a shell multiplet, prompts the computation of the integrals. As this shell driven interface matches the design of typical quantum packages, it is the most straightforward approach and requires the least modification of the existing packages.

In contrast to the use of the integral evaluators, their initialization can be very complicated, as it depends on the implementation details of the particular evaluator. The purpose of the integral evaluator factory is to hide the complexity of the initialization of integral evaluators. To create an integral evaluator, the factory must at a minimum receive the molecular basis set data. However, many integral types require the specification of additional data: derivative information, origin of the reference frame for dipole integrals, or exponents for Gaussian geminals, to name a few. The purpose of an integral descriptor (`IntegralDescriptorInterface`) is to encapsulate such data. A set of descriptors must be provided to the factory to produce corresponding integral evaluators.

Using MPQC code as an example, Figure 2 provides a code snippet outlining the procedure client codes must implement to employ CCA integral evaluators. Figure 2 assumes that the `mpqc_basis` and `services` variables have been properly assigned previously. An MPQC implementation of `MolecularInterface` is created and initialized with a pointer to the MPQC basis set object. Next, objects implementing `IntegralDescrInterface` and `DerivCentersInterface` are created (implementations of these objects are provided within the `cca-chem-generic` package).<sup>17</sup> Once a factory object is obtained from the framework, a call to the `get_evaluator4()` method, supplying the initialized descriptor and basis set objects, returns an object implementing `IntegralEvaluator4Interface`. The `compute()` method can then be used to evaluate shell quartets of integrals,

which would then be read out of `buffer` and used for calculations.

## Implementation and Benchmarking

The integral interfaces we have developed have been implemented in the Massively Parallel Quantum Chemistry (MPQC) package.<sup>27–29</sup> This is currently the only implementation of CCA integral evaluators, though NWChem<sup>30,31</sup> and GAMESS<sup>32</sup> implementations are in progress. Even this basic capability enables new functionality with the MPQC package (the native integral packages can now be mixed, see ref. 25), but the primary benefit is that now other packages may take advantage of the native MPQC integral capabilities. The performance benchmarks will also be useful for evaluation of our approach and guidance in further development. Within the native MPQC code base, two integral packages are supported. The IntV3 package is distributed with MPQC and performs integral evaluations needed for energies and gradients of conventional quantum methods. An additional package, `Cints` (based on the `Libint`<sup>33,34</sup> package), performs integral evaluations required for explicitly correlated methods in addition to those required for conventional methods. Both MPQC and `Libint` are open source, distributed under the GNU General Public License and Library GNU General Public License.

While the integral evaluation interfaces provide a framework in which to implement evaluator components, additional specifications must be imposed on the layout of data, which is passed through these interfaces. We propose CCA standards for integral data in Appendix B. The layout of integrals within buffers is of particular importance. While the `Cints` buffer layout does conform to the proposed CCA standard, the `IntV3` buffer layout does not. Most clients of integrals require a specific buffer layout, and translation costs are significant. The CCA IntV3 integral evaluator includes the option to translate integral buffers into the standard CCA format.

In general, larger amounts of time spent in MPQC's internal integral routines offset the overhead associated with the CCA interface layer, which adds several function calls and language interoperability code to each `compute()` call. High angular momentum basis sets require, on average, more work per buffer computation and therefore have decreased overhead. Gradient computations require more work per buffer than energy computations and tend to have reduced overhead. Another noticeable trend is that smaller molecules tend to have lower overhead because of the effects of integral screening. The inexpensive bound computations used in integral screening make up larger proportions of the computation as molecule size increases and more buffers are screened out. While the ultimate decision whether or not to screen out a shell multiplet lies with the integral evaluator client, our interface includes the capability for evaluators to provide the bounds information on which this decision is based. Our implementation in MPQC obtains this bounds information from the evaluators. A particular client may choose, rather, to use its own routines, eliminating a number of calls through the Babel interfaces and reducing the interface overhead somewhat for larger molecules. The test cases in the following tables are arranged to highlight these trends.

Table 1, panels A–C, reports average wall clock times for a number of energy and gradient calculations using either the native MPQC

```

using namespace Chemistry::QC::GaussianBasis;

/* already existing in MPQC */
sc::GaussianBasis mpqc_basis;
double* buffer;

/* provided by framework */
gov::cca::Services services;

/* CCA interfaces */
gov::cca::Port port;
MolecularInterface cca_basis;
IntegralDescrInterface descriptor;
DerivCentersInterface derivCenters;
IntegralEvaluatorFactoryInterface factory;
IntegralEvaluator4Interface evaluator;

/* translate basis set */
cca_basis = MPQC::GaussianBasisMolecular::_create();
cca_basis.initialize( &mpqc_basis, "basis1" );

/* create and set descriptor */
descriptor = ChemistryIntegralDescrCXX::Eri4IntegralDescr::_create();
derivCenters = ChemistryIntegralDescrCXX::DerivCenters::_create();
descriptor.set_deriv_centers( derivCenters );

/* get the factory */
port = services.getPort("IntegralEvaluatorFactoryInterface");
factory = sidl::babel_cast<IntegralEvaluatorFactoryInterface> ( port );

/* get the evaluator */
evaluator = factory.get_evaluator4( descriptor, cca_basis, cca_basis,
                                     cca_basis, cca_basis );

/* get the buffer */
buffer = evaluator.get_array(descriptor).first();

/* compute a shell quartet (written to buffer) */
evaluator.compute(0,0,0,0);

```

**Figure 2.** A C++ code snippet outlining how client code employs integral evaluators. An analogous procedure is available in each language supported by Babel (C, C++, Fortran 77, Fortran 90, Python, and Java).

integral interfaces or the CCA component interfaces. Since Cints derivative integrals have not been exposed within MPQC, these calculations were performed using the IntV3 integrals package. The IntV3 native buffer layout is used for these calculations, thus yielding a measure of the performance of the CCA integral interfaces when integral reordering is not required, the highest performance case. The CCA overheads for Hartree-Fock calculations in optimum, small molecule, big basis set cases, exemplified by the water calculations in Table 1 panel A, are under 1% and entirely insignificant. At the other end of the spectrum for Hartree-Fock methods, larger molecules with small basis sets can have significant overheads. In the benchmark suite for this study, the worst CCA overhead observed for the Hartree-Fock method was the isoprene/6-311++G\*\* energy calculation, with an overhead of 7.3%. Correlated calculations, such

as the MP2 calculations in Table 1, panel B, generally require greater computational effort subsequent to atomic orbital integral evaluation and, thus, experience moderately lower CCA overheads, which are below 5%. Density functional approaches, as surveyed in Table 1, panel C, expend significant computational effort in numerical integration, which is not currently performed through a CCA interface. Consequently, overheads for the B3LYP calculations in Table 1, panel C are uniformly low, with the greatest overhead of 3.1% seen in the water energy calculation. Extension of the evaluator interfaces to further reduce the CCA overhead is considered in our conclusions.

Table 2, panels A–C, compares average wall clock times for test cases using CCA interfaces and IntV3 integrals, comparing the use of the IntV3 native buffer layout with the CCA

**Table 1.** Average Wall Times (seconds) for (A) Hartree-Fock Calculations, (B) MP2 Calculations, and (C) B3LYP Calculations Performed Using Native and CCA Integral Interfaces.

| Test case                            | Basis set         | Native buffer layout       |                         | CCA overhead (%) |
|--------------------------------------|-------------------|----------------------------|-------------------------|------------------|
|                                      |                   | Native interface wall time | CCA interface wall time |                  |
| <b>(A) Hartree-Fock Calculations</b> |                   |                            |                         |                  |
| HF energy                            |                   |                            |                         |                  |
| Isoprene                             | 6-311++G**        | 151                        | 162                     | 7.3              |
|                                      | cc-pVDZ           | 33.5                       | 35.8                    | 6.9              |
|                                      | cc-pVTZ           | 906                        | 934                     | 3.1              |
| Phosphoserine                        | cc-pVDZ           | 254                        | 264                     | 3.9              |
| Aniline                              | cc-pVTZ           | 917                        | 942                     | 2.7              |
| Water                                | aug-cc-pV5Z       | 952                        | 955                     | 0.3              |
| HF gradient                          |                   |                            |                         |                  |
| Isoprene                             | 6-311++G**        | 252                        | 269                     | 6.7              |
|                                      | cc-pVDZ           | 92.5                       | 97.4                    | 5.3              |
|                                      | cc-pVTZ           | 1884                       | 1940                    | 3.0              |
| Phosphoserine                        | cc-pVDZ           | 652                        | 673                     | 3.2              |
| Aniline                              | cc-pVTZ           | 1795                       | 1836                    | 2.3              |
| Water                                | aug-cc-pV5Z       | 1759                       | 1768                    | 0.5              |
| <b>(B) MP2 Calculations</b>          |                   |                            |                         |                  |
| MP2 Energy                           |                   |                            |                         |                  |
| Isoprene                             | 6-311++G**        | 117                        | 120                     | 2.6              |
|                                      | cc-pVDZ           | 42.5                       | 43.3                    | 1.9              |
|                                      | cc-pVTZ           | 580                        | 590                     | 1.7              |
| Phosphoserine                        | cc-pVDZ           | 581                        | 587                     | 1.0              |
| Aniline                              | cc-pVTZ           | 566                        | 574                     | 1.4              |
| Water                                | aug-cc-pV5Z       | 259                        | 260                     | 0.4              |
| MP2 Gradient                         |                   |                            |                         |                  |
| Isoprene                             | 6-311++G**        | 689                        | 718                     | 4.2              |
|                                      | cc-pVDZ           | 236                        | 243                     | 3.0              |
|                                      | cc-pVTZ           | 4194                       | 4271                    | 1.8              |
| Phosphoserine                        | cc-pVDZ           | 2998                       | 3040                    | 1.4              |
| Aniline                              | cc-pVTZ           | 5287                       | 5370                    | 1.6              |
| Water                                | aug-cc-pV5Z       | 3845                       | 3814                    | Negligible       |
| <b>(C) B3LYP Calculations</b>        |                   |                            |                         |                  |
| B3LYP Energy                         |                   |                            |                         |                  |
| Isoprene                             | 6-311++G**        | 495                        | 508                     | 2.6              |
|                                      | cc-pVDZ           | 209                        | 210                     | 0.5              |
|                                      | cc-pVTZ           | 1853                       | 1885                    | 1.7              |
| Phosphoserine                        | 6-311G**          | 967                        | 986                     | 2.0              |
| Aniline                              | 6-311++G**        | 488                        | 495                     | 1.4              |
| Water                                | 6-311++G(3df,3pd) | 12.8                       | 13.2                    | 3.1              |
| B3LYP Gradient                       |                   |                            |                         |                  |
| Isoprene                             | 6-311++G**        | 1204                       | 1222                    | 1.5              |
|                                      | cc-pVDZ           | 722                        | 726                     | 0.6              |
|                                      | cc-pVTZ           | 4414                       | 4468                    | 1.2              |
| Phosphoserine                        | 6-311G**          | 2286                       | 2319                    | 1.4              |
| Aniline                              | 6-311++G**        | 1259                       | 1269                    | 0.8              |
| Water                                | 6-311++G(3df,3pd) | 46.6                       | 47.1                    | 1.1              |

The IntV3 integral package was used with native buffer layouts throughout. Calculations were repeated three times and averaged.

buffer layout (including buffer translation overhead). This overhead, which is in addition to the CCA interface overhead, is substantial, with Hartree-Fock reorder overheads as high as 12.3% and MP2 reorder overheads as high as 10.4%. As in the CCA

overheads, density functional approaches have moderately lower overheads for reordering, with a maximum observed overhead of 3.7%. Deeply nested loops, inefficient memory access patterns and logic to handle contractions, derivatives and differing angular types

**Table 2.** Average Wall Times (seconds) for (A) Hartree-Fock Calculations, (B) MP2 Calculations, and (C) B3LYP Calculations Performed Using CCA Integral Interfaces.

| Test case                            | Basis set         | CCA interface                  |                             | Reorder overhead (%) |
|--------------------------------------|-------------------|--------------------------------|-----------------------------|----------------------|
|                                      |                   | Native buffer layout wall time | CCA buffer layout wall time |                      |
| <b>(A) Hartree-Fock Calculations</b> |                   |                                |                             |                      |
| HF energy                            |                   |                                |                             |                      |
| Isoprene                             | 6-311++G**        | 162                            | 179                         | 10.5                 |
|                                      | cc-pVDZ           | 35.8                           | 39.4                        | 10.1                 |
|                                      | cc-pVTZ           | 934                            | 995                         | 6.5                  |
| Phosphoserine                        | cc-pVDZ           | 264                            | 280                         | 6.1                  |
| Aniline                              | cc-pVTZ           | 942                            | 991                         | 5.2                  |
| Water                                | aug-cc-pV5Z       | 955                            | 981                         | 2.7                  |
| HF gradient                          |                   |                                |                             |                      |
| Isoprene                             | 6-311++G**        | 269                            | 302                         | 12.3                 |
|                                      | cc-pVDZ           | 97.4                           | 105.2                       | 8.0                  |
|                                      | cc-pVTZ           | 1940                           | 2071                        | 6.8                  |
| Phosphoserine                        | cc-pVDZ           | 673                            | 711                         | 5.6                  |
| Aniline                              | cc-pVTZ           | 1836                           | 1945                        | 5.9                  |
| Water                                | aug-cc-pV5Z       | 1768                           | 1829                        | 3.5                  |
| <b>(B) MP2 Calculations</b>          |                   |                                |                             |                      |
| MP2 energy                           |                   |                                |                             |                      |
| Isoprene                             | 6-311++G**        | 120                            | 126                         | 5.0                  |
|                                      | cc-pVDZ           | 43.3                           | 44.3                        | 2.3                  |
|                                      | cc-pVTZ           | 580                            | 608                         | 4.8                  |
| Phosphoserine                        | cc-pVDZ           | 587                            | 598                         | 1.9                  |
| Aniline                              | cc-pVTZ           | 574                            | 585                         | 1.9                  |
| Water                                | aug-cc-pV5Z       | 260                            | 265                         | 1.9                  |
| MP2 gradient                         |                   |                                |                             |                      |
| Isoprene                             | 6-311++G**        | 718                            | 793                         | 10.4                 |
|                                      | cc-pVDZ           | 243                            | 257                         | 5.8                  |
|                                      | cc-pVTZ           | 4271                           | 4519                        | 5.8                  |
| Phosphoserine                        | cc-pVDZ           | 3040                           | 3105                        | 2.1                  |
| Aniline                              | cc-pVTZ           | 5370                           | 5633                        | 4.9                  |
| Water                                | aug-cc-pV5Z       | 3814                           | 3949                        | 3.5                  |
| <b>(C) B3LYP Calculations</b>        |                   |                                |                             |                      |
| B3LYP energy                         |                   |                                |                             |                      |
| Isoprene                             | 6-311++G**        | 508                            | 527                         | 3.7                  |
|                                      | cc-pVDZ           | 210                            | 213                         | 1.4                  |
|                                      | cc-pVTZ           | 1885                           | 1953                        | 3.6                  |
| Phosphoserine                        | 6-311G**          | 986                            | 1022                        | 3.7                  |
| Aniline                              | 6-311++G**        | 495                            | 510                         | 3.0                  |
| Water                                | 6-311++G(3df,3pd) | 13.2                           | 13.4                        | 1.5                  |
| B3LYP gradient                       |                   |                                |                             |                      |
| Isoprene                             | 6-311++G**        | 1222                           | 1257                        | 2.9                  |
|                                      | cc-pVDZ           | 726                            | 734                         | 1.1                  |
|                                      | cc-pVTZ           | 4468                           | 4615                        | 3.3                  |
| Phosphoserine                        | 6-311G**          | 2319                           | 2384                        | 2.8                  |
| Aniline                              | 6-311++G**        | 1269                           | 1296                        | 2.1                  |
| Water                                | 6-311++G(3df,3pd) | 47.1                           | 47.6                        | 1.1                  |

IntV3 integrals are used throughout, comparing native and CCA buffer orderings. Calculations were repeated three times and averaged.

are unavoidable and result in a substantially expensive reorder algorithm. Maintaining good performance with low-level interfaces requires minimizing such translation costs and, while such overhead is undoubtedly unavoidable when using legacy codes,

we strongly advocate standards adoption for any new development efforts.

Since integral packages are now interchangeable through the CCA interfaces, the source package for each of a calculation's

**Table 3.** Average Wall Times (seconds) for Calculations Performed Using CCA Integral Interfaces.

| Test case      | Basis set   | Cints wall time | Cints/IntV3 wall time | IntV3 overhead (%) |
|----------------|-------------|-----------------|-----------------------|--------------------|
| MP2-R12 energy |             |                 |                       |                    |
| Isoprene       | cc-pVDZ     | 246.2           | 246.2                 | 0.0                |
| Water          | aug-cc-pV5Z | 1576            | 1577                  | 0.1                |

Calculations using Cints integrals throughout and using Cints two-electron integrals along with IntV3 one-electron integrals are compared. CCA buffer layouts (identical to Cints buffer layouts) were used throughout. Calculations were repeated three times and averaged.

integral types can be selected individually. New packages implementing advanced integral types no longer need to duplicate standard integral capabilities, as these are available from a growing number of CCA implementations based on mature integral packages. Table 3 provides timings for two MP2-R12/A' test cases demonstrating this flexibility. For the first set of calculations, Cints integrals were used throughout. For comparison, the second set of calculations use Cints for two-electron integrals while IntV3 is used to provide overlap and core Hamiltonian integrals. Though reordering of the IntV3 buffers was required for the second set of calculations, very little overhead is observed since the cost of one-electron integrals is minor. These calculations also highlight opportunities for "quality of service" improvements, which the CCA architecture enables. With a large set of interchangeable integral evaluator implementations, it is possible to include components for automated selection of the most appropriate implementation for a given integral type, calculation type, and hardware environment.

### Embedded Frameworks

The components previously developed for geometry optimization<sup>16,17</sup> encapsulated the high-level functionalities of various domain-specific packages. There was no existing support for constructing and configuring applications based on the disparate packages, making an entirely component-based architecture composed and configured via the framework interface a natural choice. While direct interaction with the framework allows tremendous flexibility, it is unlikely that end-users will have the skills necessary to properly configure component applications. The development of a standard input format that generically supports quantum chemistry packages was recognized as both a unique opportunity and a daunting challenge of this design, and has not, thus far, been attempted.

Shifting some low-level functionality to components more fully realizes the potential of component technology to facilitate interoperable, rather than just interchangeable, chemistry packages, and suggests the possibility of applications which only utilize components for specific tasks. The CCA specification and Ccaffeine framework do support embedding of framework functionality in stand-alone codes, allowing mixed legacy/component application architectures. Surely, each chemistry package supports configuration of computations, and this configuration capability can be easily extended to support accessing low-level functionality provided by

components. Embedding component composition and configuration inside legacy codes allows package developers to expose to end-users only those features useful in a particular context. The barrier for end-users to adopt component technology is substantially reduced; the stand-alone package is run as before with a small number of extra parameters introduced to calculation inputs. An embedded framework environment has been added along with the integral component client and server code that we have added to the MPQC package, allowing MPQC chemistry models to access integrals via components in both exclusively component and mixed native/component modes. The adoption of integral components by MPQC users is now trivial.

### Conclusions

Through the design and implementation of an extendable interface for molecular integral evaluation in quantum chemistry, we have demonstrated the efficacy of exposing low-level software capabilities using component approaches. For integral evaluation, calls through component interfaces easily number in the millions and involve several extra function calls and a fair amount of language interoperability code, yet overheads are reasonable. For calculations on small molecules using large basis sets, interface overheads below 1% can be expected. For larger molecules and smaller basis sets, interface overheads are clearly more significant but are acceptable, with all benchmark values falling within 10%. While overheads in the 10% range are likely large enough to discourage the use of these interfaces when high performance native implementations are available, the ability to share integral types and rapidly implement new approaches ensures the usefulness of integral components. As has been demonstrated for mesh interfaces,<sup>13</sup> decreasing the granularity of component interactions decreases overheads. Though it would require a much larger modification of the existing quantum packages, a vector driven approach, where multiple shell multiplets are computed for each call through the CCA interface, could substantially reduce the CCA overhead as well as facilitate more effective use of the single-instruction multiple-data (vector) hardware now ubiquitous in general purpose CPU's. Moving forward, we plan to implement such a vector interface.

When low-level software functionality is shared between packages, adapting arbitrary implementation details, such as integral buffer ordering, to a common standard can cause substantial overhead, which overshadows the costs of an added interface layer. In the case of integral buffer ordering, our benchmarks show overheads as high as 12%, in addition to the interface overhead due to the CCA layer. While such significant overheads will often be unavoidable when adapting legacy codes for component implementations, the development and adoption of standards in new work is critical to reducing such inefficiencies in the long term, allowing for a high performance community code base. We set forth standards for integral evaluation in Appendix B and encourage their adoption.

As our efforts in developing a chemistry component toolkit have progressed, it has become evident that the CCA approach is about much more than components. The true strength of the CCA is as a comprehensive approach to scientific software engineering. Holding as much or more importance than component concepts are the approaches to interface and data standards, language and package

interoperability, and collaborative, community-based development, which have naturally evolved along with component standards and middleware. In our view, the component concept has functioned as a catalyst, which has focused a large community of computational and computer scientists on solving the problems of large-scale, collaborative, scientific software development. While improving usability and maintenance of this increasingly complex chemistry software project remains a challenge, this work is significant progress in developing an interoperable code base for quantum chemistry.

## Calculation Details

The codes implementing CCA integral interfaces are currently under development and will be available in forthcoming releases. Source code for the cca-chem-generic package,<sup>17</sup> which provides interface definitions and some implementations that are generally useful, was a snapshot of the babel-1-0-branch CVS branch as of 2/21/2007. MPQC<sup>27-29</sup> source code was a snapshot of MPQC's babel-1-0-branch CVS branch as of 2/15/2007. Babel<sup>11,12</sup> source code was a snapshot of Babel's 1.1 development Subversion branch as of 12/21/2006. Source code for the cca-tools package<sup>5</sup> was the 0.6.1\_rc2 release, with slight modifications to allow building against Babel 1.1. Ccaffeine,<sup>6</sup> one of several frameworks<sup>6-10</sup> that comply with the CCA specification, was used. All codes were built using the gcc 3.4.3 compiler with the default x86\_64 processor target and -O2 optimization. Benchmarks were performed in single process, single thread mode on Intel Xeon 5160 CPU's (Core 2 / Woodcrest architecture) clocked at 3.00 GHz, running Red Hat Enterprise Linux AS release 4.

## Acknowledgements

The authors acknowledge the contributions of Dr. David E. Bernholdt and Dr. Ricky A. Kendall during the initial interface design. We also thank Dr. David E. Bernholdt for valuable comments on an early draft of this paper. This work has been supported in part by the U.S. Department of Energy's Scientific Discovery through Advanced Computing (SciDAC) initiative,<sup>35</sup> through the Center for Component Technology for Terascale Simulation Software, of which Oak Ridge, Pacific Northwest, and Sandia National Laboratories are members. This research was performed in part using the Molecular Science Computing Facility (MSCF) in the William R. Wiley Environmental Laboratory at the Pacific Northwest National Laboratory (PNNL). The MSCF is funded by the Office of Biological and Environmental Research in the U.S. Department of Energy. PNNL is operated by Battelle for the U.S. Department of Energy. Oak Ridge National Laboratory is managed by UT-Battelle, LLC, for the U.S. Department of Energy. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the U.S. Department of Energy.

## Appendix A: Integral Evaluation Interfaces

Interfaces concerned with integral evaluation over Gaussian basis functions are described in the following sections. Babel's Scientific Interface Description Language (SIDL)<sup>11</sup> is a neutral choice

for this description. The SIDL code for the following interfaces resides in the Chemistry.QC.GaussianBasis package, a subsection of the chemistry.sidl file distributed with the cca-chem-generic software.<sup>17</sup>

For convenience, an enumeration for function angular types is provided.

```
enum AngularType { CARTESIAN, SPHERICAL, MIXED }
```

Codes which utilize integral evaluator components must pass Gaussian basis set data to the evaluators. This task is performed by copying basis set data into class implementations of the ShellInterface, AtomicInterface, and MolecularInterface and passing the MolecularInterface object to the integral evaluators (a MolecularInterface object is any SIDL class that implements the MolecularInterface). A MolecularInterface object, representing a molecular basis set, contains a set of AtomicInterface objects, each of which contains a set of ShellInterface objects. The MolecularInterface object also contains a MoleculeInterface object, which provides basic molecule properties, such as the geometry.

### ShellInterface

```
int get_n_contraction()
```

*Get the number of contractions in the shell.*

**Returns:**

*number of contractions*

```
int get_n_primitive()
```

*Get the number of primitives in the shell.*

**Returns:**

*number of primitives*

```
double get_contraction_coef(in int connum, in int expnum)
```

*Get the coefficient for an unnormalized primitive.*

**Returns:**

*contraction coefficient*

**Parameters:**

**connum** *contraction number*

**expnum** *primitive number*

```
double get_exponent(in int expnum)
```

*Get the exponent for a primitive.*

**Returns:**

*exponent*

**Parameters:**

**expnum** *primitive number*

```
int get_angular_momentum(in int connum)
```

*Get the angular momentum for a single contraction.*

**Returns:**

*angular momentum value*

**Parameters:**

**connum** *contraction number*



**int get\_max\_angular\_momentum()**

Get the max angular momentum, considering all contractions in the shell.

**Returns:**

maximum angular momentum value

**AngularType get\_contraction\_angular\_type(in int connum)**

Get the angular type for a single contraction.

**Returns:**

angular type

**Parameters:**

**connum** contraction number

**AngularType get\_angular\_type()**

Get the shell angular type.

**Returns:**

angular type

**void print\_shell()**

Print the shell data.

**AtomicInterface****string get\_name()**

Get the canonical basis set name.

**Returns:**

canonical name

**int get\_n\_basis()**

Get the number of basis functions.

**Returns:**

number of functions

**int get\_n\_shell()**

Get the number of shells.

**Returns:**

number of shells

**int get\_max\_angular\_momentum()**

Get the max angular momentum for any shell on the atom.

**Returns:**

max angular momentum value

**AngularType get\_angular\_type()**

Get the angular type for the atom.

**Returns:**

angular type

**ShellInterface get\_shell(in int shellnum)**

Get a gaussian shell.

**Returns:**

object implementing **ShellInterface**

**Parameters:**

**shellnum** shell number

**void print\_atomic()**

Print the atomic basis data.

**MolecularInterface****string get\_label()**

Get the user specified name.

**Returns:**

name

**long get\_n\_basis()**

Get the number of basis functions.

**Returns:**

number of functions

**long get\_n\_shell()**

Get the number of shells.

**Returns:**

number of shells

**int get\_max\_angular\_momentum()**

Get the max angular momentum for any contraction in the basis set.

**Returns:**

max angular momentum value

**AngularType get\_angular\_type()**

Get the angular type.

**Returns:**

angular type

**AtomicInterface get\_atomic(in long atomnum)**

Get an atomic basis set.

**Returns:**

object implementing the **AtomicInterface**

**Parameters:**

**atomnum** atom number

**MoleculeInterface get\_molecule()**

Get the molecule.

**Returns:**

object implementing the **MoleculeInterface**

**void print\_molecular()**

Print the molecular basis data.

**MoleculeInterface****void initialize(in long natom, in string unitname)**

Initialize a molecule.

**Parameters:**

**natom** number of atoms

**unitname** units for coordinates

**Physics.UnitsInterface get\_units()**

Returns a units object that corresponds to the units that are used by `get_cart_coord()` or `set_cart_coord()`.

**Returns:**

units object

**long get\_n\_atom()**

Returns the number of atoms.

**Returns:**

*number of atoms*

**int get\_atomic\_number(in long atomnum)**

Returns the atomic number of an atom.

**Returns:**

*atomic number*

**Parameters:**

**atomnum** atom index

**void set\_atomic\_number(in long atomnum, in int atomic\_number)**

Sets atomic number of an atom.

**Parameters:**

**atomnum** atom id

**atomic\_number** atomic number

**double get\_net\_charge()**

Returns net charge of the molecule.

**Returns:**

*net charge*

**void set\_net\_charge(in double charge)**

Sets the net charge of the molecule.

**Parameters:**

**charge** molecular charge

**double get\_charge(in long atomnum)**

Returns integer charge at an atom.

**Returns:**

*integer charge at an atom*

**Parameters:**

**atomnum** atom index

**void set\_charge(in long atomnum, in double charge)**

Sets charge at an atom.

**Parameters:**

**atomnum** atom index

**charge** charge at an atom

**array<double,1> get\_coor()**

Returns the Cartesian coordinate array.

**Returns:**

*Cartesian coordinate array*

**void set\_coor(in array<double,1> x)**

Sets the Cartesian coordinates.

**Parameters:**

**x** coordinate array

**double get\_cart\_coor(in long atomnum, in int xyz)**

**Returns:**

*a Cartesian coordinate*

**Parameters:**

**atomnum** atom index

**xyz** give 0 for x, 1 for y, and 2 for z

**void set\_cart\_coor(in long atomnum, in int xyz, in double val)**

Sets individual Cartesian coordinate.

**Parameters:**

**atomnum** atom index

**xyz** give 0 for x, 1 for y, and 2 for z

**val** coordinate value

**string get\_atomic\_label(in long atomnum)**

Returns programmer-defined label for an atom.

**Returns:**

*label*

*programmer-defined label*

**Parameters:**

**atomnum** atom index

**void set\_atomic\_label(in long atomnum, in string label)**

Sets programmer-defined label for an atom.

**Parameters:**

**atomnum** atom index

**label** programmer-defined label

A MoleculeInterface object contains a Physics.UnitsInterface object which provides units for coordinate values.

## Physics.UnitsInterface

**void initialize( in string unitname )**

Initializes the units as a string.

**Parameters:**

**unitname** "angstroms" or "bohr"

**string get\_unit\_name()**

Returns the units as a human readable string.

**Returns:**

*unit name*

**double convert\_to(in string unitname)**

Returns conversion factor from self's units to the given unit name.

**Returns:**

*conversion factor*

**Parameters:**

**unitname** unit name

**double convert\_from(in string unitname)**

Returns conversion factor to self's units from the given unit name.

**Returns:**

*conversion factor*

**Parameters:**

**unitname** unit name

Many integral types require data in addition to integral buffers to be shared between servers and clients. A number of data container utility interfaces serve this function, and this collection will grow as client/server components with additional capabilities are implemented. The first such utility interface we describe is the `DerivCentersInterface`, used to specify details for nuclear derivative calculations. Code utilizing an integral evaluator must pass a `DerivCentersInterface` object to the evaluator. For derivatives with operators which are independent of nuclear coordinates, translational invariance allows derivatives with respect to one center to be omitted, and details about omitted centers are shared through this interface. For derivatives of operators which are dependent upon nuclear coordinates, derivatives with respect to every atom must be taken, and the derivative atom information is shared using the `set_deriv_atom()` and `get_deriv_atom()` methods. The *segment* number, defined as the number of basic buffer segments a buffer is composed of, is also shared through this interface. As an example of the buffer segment concept, a dipole integral buffer consists of one buffer segment for each of *x*, *y*, and *z*, yielding a segment number of 3.

### DerivCentersInterface

#### **void clear()**

Clear the list of centers.

#### **Returns:**

*list of centers*

#### **void add\_center(in long center, in long atom)**

Add a center for which derivatives will be computed.

#### **Parameters:**

**center** center number (between 0 and 3 inclusive)

**atom** atom number corresponding to center

#### **void add\_omitted(in long center, in long atom)**

Add a center for which derivatives will not be computed.

#### **Parameters:**

**center** center number (between 0 and 3 inclusive)

**atom** atom number corresponding to center

#### **long n()**

Returns the number of centers for which derivatives will be computed.

#### **Returns:**

*number of centers*

#### **long center(in long i)**

Returns center number.

#### **Returns:**

*center number (between 0 and 3 inclusive)*

#### **Parameters:**

**i** computed center index (between 0 and *n()*-1 inclusive)

#### **long atom(in long i)**

Returns atom number.

#### **Returns:**

*atom number*

#### **Parameters:**

**i** computed center index (between 0 and *n()*-1 inclusive)

#### **long omitted\_center()**

Returns the omitted center number.

#### **Returns:**

*omitted center number*

#### **int has\_omitted\_center()**

Returns 1 if there is an omitted center.

#### **Returns:**

*1 (true) or 0 (false)*

#### **long omitted\_atom()**

Returns atom that is omitted from the integral buffer.

#### **Returns:**

*omitted atom number*

#### **void set\_deriv\_atom(in int deriv\_atom)**

Set the atom a derivative is taken with respect to.

#### **Parameters:**

**deriv\_atom** atom number

#### **int get\_deriv\_atom()**

Get the atom a derivative is taken with respect to.

#### **Returns:**

*atom number*

Currently, the only additional data container utility interface is the `DipoleDataInterface` which provides origin information for dipole and quadrupole integrals.

### DipoleDataInterface

#### **void set\_origin( in array<double> origin )**

Set the dipole origin.

#### **Parameters:**

**origin** Cartesian coordinate array

#### **array<double> get\_origin()**

Get the dipole origin.

#### **Returns:**

*Cartesian coordinate array*

When an integral evaluator is requested, a composite of `IntegralDescrInterface` objects for the requested integral types is passed to the evaluator factory. This action provides both the list of required integral types and any additional data required. The object oriented features of Babel allow a collection of derived integral descriptors to be upcast to a collection of base descriptors which are passed through the `EvaluatorFactoryInterface`. The evaluator factory then checks descriptor types, performs any necessary downcasts, and thereby obtains the auxiliary data. The `EvaluatorFactoryInterface` is thus generic and extendable for all possible integral types, requiring only the implementation of a new derived integral descriptor for types requiring additional data. We now describe the `IntegralDescrInterface`, `CompositeIntegralDescrInterface`, and

IntegralEvaluatorFactoryInterface, along with an example derived descriptor interface, the DipoleIntegralDescrInterface.

### IntegralDescrInterface

**string get\_type()**

*Get integral type.*

**Returns:**  
*type*

**int get\_n\_segment()**

*Get number of segments.*

**Returns:**  
*number of segments*

**void set\_deriv\_lvl( in int deriv )**

*Set derivative level.*

**Parameters:**  
*deriv derivative level*

**int get\_deriv\_lvl()**

*Get derivative level.*

**Returns:**  
*derivative level*

**void set\_deriv\_centers( in DerivCentersInterface dc )**

*Set derivative centers object.*

**Parameters:**  
*dc derivative centers object*

**DerivCentersInterface get\_deriv\_centers()**

*Get derivative centers object.*

**Returns:**  
*derivative centers object*

### DipoleIntegralDescrInterface

extends **IntegralDescrInterface**

**void set\_dipole\_data( in DipoleDataInterface dipole\_data )**

*Set the dipole data.*

**Parameters:**  
*dipole\_data dipole data*

**DipoleDataInterface get\_dipole\_data()**

*Get the dipole data.*

**Returns:**  
*dipole data*

### CompositeIntegralDescrInterface

**void add\_descr( in IntegralDescrInterface desc )**

*Add an integral descriptor.*

**Parameters:**  
*desc integral descriptor*

**int get\_n\_descr()**

*Get number of descriptors contained.*

**Returns:**  
*number of descriptors*

**IntegralDescrInterface get\_descr( in int n )**

*Get an integral descriptor.*

**Returns:**  
*integral descriptor*  
**Parameters:**  
*n descriptor index*

**int is\_contained( in IntegralDescrInterface desc )**

*Query if a descriptor with matching type and derivative level is contained.*

**Returns:**  
*1 (true) or 0 (false)*

**Parameters:**  
*desc integral descriptor*

**void clear()**

*Clear all descriptors.*

### IntegralEvaluatorFactoryInterface

**string get\_name()**

*Get factory name.*

**Returns:**  
*name*

**CompositeIntegralDescrInterface get\_descriptor()**

*Get composite of descriptors for supported integrals.*

**Returns:**  
*composite integral descriptor*

**bool is\_supported( in IntegralDescrInterface desc )**

*Query if a type and derivative level is supported.*

**Returns:**  
*true or false*

**Parameters:**  
*desc integral descriptor*

**void set\_storage( in long storage )**

*Set storage that the factory is allowed to utilize.*

**Parameters:**  
*storage allowed storage in bytes*

### IntegralEvaluator1Interface

**get\_evaluator1(**

**in CompositeIntegralDescrInterface desc,**  
**in MolecularInterface bs1**

**)**

*Get a 1-center integral evaluator.*

**Returns:**  
*1-center evaluator*

**Parameters:**  
*desc composite integral descriptor*  
*bs1 basis set for center 1*

**IntegralEvaluator2Interface**

```
get_evaluator2(
    in CompositeIntegralDescrInterface desc,
    in MolecularInterface bs1,
    in MolecularInterface bs2
)
```

*Get a 2-center integral evaluator.*

**Returns:**

*2-center evaluator*

**Parameters:**

**desc** *composite integral descriptor*

**bs1** *basis set for center 1*

**bs2** *basis set for center 2*

**IntegralEvaluator3Interface**

```
get_evaluator3(
    in CompositeIntegralDescrInterface desc,
    in MolecularInterface bs1,
    in MolecularInterface bs2,
    in MolecularInterface bs3
)
```

*Get a 3-center integral evaluator.*

**Returns:**

*3-center evaluator*

**Parameters:**

**desc** *composite integral descriptor*

**bs1** *basis set for center 1*

**bs2** *basis set for center 2*

**bs3** *basis set for center 3*

**IntegralEvaluator4Interface**

```
get_evaluator4(
    in CompositeIntegralDescrInterface desc,
    in MolecularInterface bs1,
    in MolecularInterface bs2,
    in MolecularInterface bs3,
    in MolecularInterface bs4
)
```

*Get a 4-center integral evaluator.*

**Returns:**

*4-center evaluator*

**Parameters:**

**desc** *composite integral descriptor*

**bs1** *basis set for center 1*

**bs2** *basis set for center 2*

**bs3** *basis set for center 3*

**bs4** *basis set for center 4*

**int finalize()**

*This should be called when the object is no longer needed.*

*No other members may be called after finalize.*

**Returns:**

*0 on success*

evaluator factories (following the Abstract Factory Pattern<sup>36</sup>). Once client code provides rules for the routing of integral evaluator requests, the super factory acts as a single evaluator factory enveloping the capabilities of all connected factories.

**IntegralSuperFactoryInterface**

extends **IntegralEvaluatorFactoryInterface**

```
array<string,1> add_uses_ports(in int n)
```

*Add uses ports to component implementation.*

**Returns:**

*array of uses port names*

**Parameters:**

**n** *number of additional uses ports*

```
void remove_port(in int portid)
```

*Remove uses port.*

**Parameters:**

**portid** *port index*

```
array<string,1> get_port_names()
```

*Get uses port names.*

**Returns:**

*array of uses port names*

```
array<string,1> get_factory_names()
```

*Get attached evaluator factory names.*

**Returns:**

*array of factory names*

```
array<CompositeIntegralDescrInterface,1> get_descriptors()
```

*Get composite of available integral descriptors for each factory.*

**Returns:**

*array of composite integral descriptors*

```
void set_default_subfactory(in string fac )
```

*Set the default factory for integral evaluator requests.*

**Parameters:**

**fac** *name of default factory*

```
void
```

```
set_subfactory_config(
```

```
    in array<string,1> types,
```

```
    in array<string,1> derivs,
```

```
    in array<string,1> facs
```

```
)
```

*Configure which factory handles specific integral type requests.*

**Parameters:**

**types** *array of integral types*

**derivs** *array of derivative levels (integer or "n" for wildcard)*

**wildcard**

**facs** *array of factory names*

The **IntegralSuperFactoryInterface** provides a management layer for simplifying the use of multiple integral

The remaining interfaces specify the integral evaluator interfaces themselves. A base interface, **IntegralEvaluator**

Interface, is extended for one, two, three, and four-center integrals. We describe the two-center interface here; extension to other numbers of centers is obvious.

### IntegralEvaluatorInterface

**array<double> get\_array(in IntegralDescrInterface desc)**

Get sidl array buffer for given type.

**Returns:**

*sidl array*

**Parameters:**

**desc** *integral descriptor*

**CompositeIntegralDescrInterface get\_descriptor()**

Get composite of descriptors for integral types supported.

**Returns:**

*composite integral descriptor*

**int finalize()**

This should be called when the object is no longer needed.

No other members may be called after finalize.

**Returns:**

*0 on success*

### IntegralEvaluator2Interface

**void compute(in long shellnum1, in long shellnum2)**

Compute all buffers for specified shell multiplet.

**Parameters:**

**shellnum1** *shell 1 index*

**shellnum2** *shell 2 index*

**double compute\_bounds(in long shellnum1, in long shellnum2)**

Compute max integral bound.

**Returns:**

*max integral bound for all computed types*

**Parameters:**

**shellnum1** *shell 1 index*

**shellnum2** *shell 2 index*

**array<double> compute\_bounds\_array(in long shellnum1, in long shellnum2);**

Compute integral bounds for each computed type.

**Returns:**

*SIDL array of integral bounds*

**Parameters:**

**shellnum1** *shell 1 index*

**shellnum2** *shell 2 index*

## Appendix B: Proposed Integral Standards

While the integral interfaces we have proposed define a set of function calls which may be used to obtain, initialize, and utilize molecular integral evaluators, standards must be specified for

implementation details, namely buffer layout and normalization conventions.

### Buffer Layout

The most intuitive algorithm for the ordering of Cartesian functions is proposed. Given angular momentum  $l$ , the Cartesian functions  $x^a y^b z^c$  are ordered as follows

starting with

$$a = l$$

$$b = c = 0$$

the next function is given by

$$\text{if } (c < l - a) \{ \\ b = b - 1 \\ c = c + 1 \\ \}$$

else {

$$a = a - 1$$

$$c = 0$$

$$b = l - a$$

}

For example, a  $d$  shell is ordered

$$x^2, xy, xz, y^2, yz, z^2$$

For indexing within a Cartesian shell multiplet buffer, the first center is treated as the most significant, with each subsequent center receiving less significance.

For a  $pp$  shell doublet the ordering is

$$(x|x) (x|y) (x|z)$$

$$(y|x) (y|y) (y|z)$$

$$(z|x) (z|y) (z|z)$$

For an  $sppp$  shell quartet the ordering is

$$(1x|xx) (1x|xy) (1x|xz)$$

$$(1x|yx) (1x|yy) (1x|yz)$$

$$(1x|zx) (1x|zy) (1x|zz)$$

$$(1y|xx) (1y|xy) (1y|xz)$$

$$(1y|yx) (1y|yy) (1y|yz)$$

$$(1y|zx) (1y|zy) (1y|zz)$$

$$(1z|xx) (1z|xy) (1z|xz)$$

$$(1z|yx) (1z|yy) (1z|yz)$$

$$(1z|zx) (1z|zy) (1z|zz)$$

Note that redundant integrals may be included. The ordering within a pure angular momentum buffer follows the same significance rule, with functions ordered in decreasing  $m_l$  ( $l, l - 1, \dots, -l$ ).

For an  $n$ -center multiplet, a first derivative buffer contains a set of three derivative multiplets ( $\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z}$ ) for each of up to  $n - 1$  unique centers (at least one center omitted due to translational invariance).

For an *sssp* first derivative shell quartet (omitting derivatives with respect to center four) the ordering is

$$\begin{aligned} & \frac{\partial}{\partial x_1} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial}{\partial y_1} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial}{\partial z_1} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial}{\partial x_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial}{\partial y_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial}{\partial z_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial}{\partial x_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial}{\partial y_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial}{\partial z_3} [(11|1x) \ (11|1y) \ (11|1z)] \end{aligned}$$

Similarly, for second derivatives the ordering is

$$\begin{aligned} & \frac{\partial^2}{\partial x_1^2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_1 \partial y_1} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_1 \partial z_1} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_1 \partial x_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_1 \partial y_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_1 \partial z_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_1 \partial x_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_1 \partial y_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_1 \partial z_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_1^2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_1 \partial z_1} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_1 \partial x_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_1 \partial y_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_1 \partial z_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_1 \partial x_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_1 \partial y_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_1 \partial z_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial z_1^2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial z_1 \partial x_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial z_1 \partial y_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial z_1 \partial z_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial z_1 \partial x_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial z_1 \partial y_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial z_1 \partial z_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_2^2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_2 \partial y_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_2 \partial z_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_2 \partial x_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_2 \partial y_3} [(11|1x) \ (11|1y) \ (11|1z)] \end{aligned}$$

$$\begin{aligned} & \frac{\partial^2}{\partial x_2 \partial z_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_2^2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_2 \partial z_2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_2 \partial x_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_2 \partial y_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_2 \partial z_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial z_2^2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial z_2 \partial x_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial z_2 \partial y_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial z_2 \partial z_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_3^2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_3 \partial y_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial x_3 \partial z_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_3^2} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial y_3 \partial z_3} [(11|1x) \ (11|1y) \ (11|1z)] \\ & \frac{\partial^2}{\partial z_3^2} [(11|1x) \ (11|1y) \ (11|1z)] \end{aligned}$$

and likewise for higher order derivatives.

### Normalization

Gaussian integral packages can have significantly different normalization conventions for the target integrals. The convention often depends on implementation details, such as the evaluation method. The normalization convention we adopted is “natural” for the majority of integral packages aimed at basis sets with segmented contractions.

Cartesian Gaussian functions in a shell of angular momentum  $L$  have the same normalization factor  $N$ .  $N$  is determined such that the Cartesian functions  $x^L$ ,  $y^L$ , and  $z^L$  are normalized to unity. The norm of a Cartesian Gaussian  $x^a y^b z^c$  is therefore

$$\|x^a y^b z^c\| = \frac{(2a + 2b + 2c - 1)!!}{(2a - 1)!!(2b - 1)!!(2c - 1)!!}$$

All spherical harmonic Gaussians are normalized to unity. Transformation from Cartesian to spherical harmonics Gaussians was described in detail by Schlegel and Frisch.<sup>37</sup>

### References

1. van der Steen, A. J. Overview of Recent Supercomputers; Available at <http://www.phys.uu.nl/~euroben/reports/web05a/overview.html>.
2. Földesy, P. In Proceedings of the 2004 IEEE International Joint Conference on Neural Networks; IEEE Press: Santa Fe, 2004.
3. Armstrong, R.; Gannon, D.; Geist, A.; Keahey, K.; Kohn, S.; McInnes, L.; Parker, S.; Smolinski, B. In Proceedings of the Eighth International Symposium on High Performance Distributed Computing; IEEE Press: Santa Fe, 1999.

4. Bernholdt, D. E.; Allan, B. A.; Armstrong, R.; Bertrand, F.; Chiu, K.; Dahlgren, T. L.; Damevski, K.; Elwasif, W. R.; Epperly, T. G. W.; Govindaraju, M.; Katz, D. S.; Kohl, J. A.; Krishnan, M.; Kumpfert, G.; Larson, J. W.; Lefantzi, S.; Lewis, M. J.; Malony, A. D.; McInnes, L. C.; Nieplocha, J.; Norris, B.; Parker, S. G.; Ray, J.; Shende, S.; Windus, T. L.; Zhou, S. *Intl J High-Perf Comput Appl* 2006, 20, 163.
5. CCA Forum; CCA Forum homepage; Available at <http://www.cca-forum.org/>.
6. Allan, B. A.; Armstrong, R. C.; Wolfe, A. P.; Ray, J.; Bernholdt, D. E.; Kohl, J. A. *Concurrency Comput Pract Exp* 2002, 14, 1.
7. Govindaraju, M.; Krishnan, S.; Chiu, K.; Slominski, A.; Gannon, D.; Bramley, R. In 3rd IEEE/ACM International Symposium on Cluster Computing and the Grid; IEEE Press: Santa Fe, 2003.
8. Zhang, K.; Damevski, K.; Venkatachalapathy, V.; Parker, S. In Proceedings of the 9th International Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS 2004); IEEE Press: Santa Fe, pp. 72–79, 2004.
9. Grimshaw, A.; Ferrari, A.; Knabe, F.; Humphrey, M. *Computer* 1999, 32, 29.
10. Lewis, M.; Ferrari, A.; Humphrey, M.; Karpovich, J.; Morgan, M.; Natrajan, A.; Nguyen-Tuong, A.; Wasson, G.; Grimshaw, A. *J Parallel Distribut Comp* 2003, 63, 525.
11. Dahlgren, T.; Epperly, T.; Kumpfert, G. *Babel User's Guide*, version 0.9.0, 2004.
12. Lawrence Livermore National Laboratory; Babel homepage; <http://www.llnl.gov/CASC/components/babel.html>.
13. McInnes, L. C.; Allan, B. A.; Armstrong, R.; Benson, S. J.; Bernholdt, D. E.; Dahlgren, T. L.; Diachin, L. F.; Krishnan, M.; Kohl, J. A.; Larson, J. W.; Lefantzi, S.; Nieplocha, J.; Norris, B.; Parker, S. G.; Ray, J.; Zhou, S. In *Solutions of PDEs on Parallel Computers*; Springer-Verlag: New York, 2005; Chapter Parallel PDE-Based Simulations Using the Common Component Architecture.
14. Lefantzi, S.; Ray, J.; Najm, H. N. In Proceedings of the 17th International Parallel and Distributed Processing Symposium (IPDPS 2003); IEEE Press: Santa Fe, 2003.
15. Lefantzi, S.; Ray, J.; Kennedy, C. A.; Najm, H. N. *Prog Comput Fluid Dyn* 2005, 5, 298.
16. Kenny, J. P.; Benson, S. J.; Alexeev, Y.; Sarich, J.; Janssen, C. L.; McInnes, L. C.; Krishnan, M.; Nieplocha, J.; Jurrus, E.; Fahlstrom, C.; Windus, T. L. *J Comput Chem* 2004, 25, 1717.
17. CCA Quantum Chemistry Working Group; Homepage; Available at <http://www.cca-forum.org~cca-chem>.
18. Krishnan, M.; Alexeev, Y.; Windus, T. L.; Nieplocha, J. In Proceedings of the 2005 ACM/IEEE Conference on Supercomputing; IEEE Computer Society: Washington, DC 2005.
19. Klopper, W. In *Encyclopedia of Computational Chemistry*; v. R. Schleyer, P.; Allinger, N. L.; Clark, T.; Gasteiger, J.; Kollman, P. A.; Schaefer III, H. F.; Scheiner, P. R., Eds.; Wiley: Chichester, UK, 1998.
20. Rychlewski, J., Ed. *Progress in Theoretical Chemistry and Physics*, Vol. 13, Springer: New York, 2004.
21. Klopper, W.; Manby, F. R.; Ten-no, S.; Valeev, E. F. *Int Rev Phys Chem* 2006, 25, 427.
22. Douglas, M.; Kroll, N. M. *Ann Phys* 1974, 82, 89.
23. Hess, B. A. *Phys Rev A* 1986, 33, 3742.
24. Pacios, I. F.; Christiansen, P. A. *J Chem Phys* 1985, 82, 2664.
25. Janssen, C. L.; Kenny, J. P.; Nielsen, I. M. B.; Krishnan, M.; Guru-moorthi, V.; Valeev, E. F.; Windus, T. L. *J Phys: Conf Ser* 2006, 46, 220.
26. Szyperski, C. *Component Software: Beyond Object-Oriented Programming*; Addison-Wesley: New York, 2002; Chapter 4.
27. Janssen, C. L.; Nielsen, I. M. B.; Colvin, M. E. *Encyclopedia of Computational Chemistry*; Wiley: Chichester, UK, 1998; Chapter Parallel Processing for ab Initio Quantum Mechanical Methods.
28. Sandia National Laboratories; MPQC homepage; Available at <http://www.mpqc.org/>.
29. Janssen, C.; Seidl, E.; Colvin, M. In *ACS Symposium Series, Parallel Computing in Computational Chemistry*, Vol. 592; American Chemical Society: Washington, DC 1995.
30. Kendall, R. A.; Apra, E.; Bernholdt, D. E.; Bylaska, E. J.; Dupuis, M.; Fann, G. I.; Harrison, R. J.; Ju, J. L.; Nichols, J. A.; Nieplocha, J.; Straatsma, T. P.; Windus, T. L.; Wong, A. T. *Comput Phys Commun* 2000, 128, 260.
31. Pacific Northwest National Laboratory; NWChem homepage; Available at <http://www.emsl.pnl.gov/docs/nwchem/>.
32. Gordon research group, Iowa State University; GAMESS homepage; Available at <http://www.msg.ameslab.gov/GAMESS/>.
33. Virginia Tech; Valeev, E. F. Libint library; Available at <http://www.chem.vt.edu/chem-dept/valeev/libint/>.
34. Valeev, E. F.; Janssen, C. L. *J Chem Phys* 2004, 121, 1214.
35. U.S. Dept. of Energy; SciDAC Initiative homepage; Available at <http://www.osti.gov/scidac/>.
36. Gamma, E.; Helm, R.; Johnson, R.; Vliissides, J. *Design Patterns, Elements of Reusable Object-Oriented Software*; Addison-Wesley: New York, 1998; Chapter Creational Patterns.
37. Schlegel, H. B.; Frisch, M. J. *Int J Quant Chem* 1995, 54, 83.