

PXFS: A Persistent Storage Model for Extreme Scale

Shuangyang Yang^{*§}, Maciej Brodowicz[†], Walter B Ligon III[‡], Hartmut Kaiser^{*§}

^{*}Center for Computation and Technology, Louisiana State University, Baton Rouge, LA

[†]Center for Research in Extreme Scale Technology, Indiana University, Bloomington, IN

[‡]Parallel Architecture Research Laboratory, Clemson University, Clemson, SC

[§]Department of Computer Science, Louisiana State University, Baton Rouge, LA
syang16@lsu.edu, mbrodowi@indiana.edu, walt@clemson.edu, hkaiser@cct.lsu.edu

Abstract—The continuing technological progress resulted in sustained increase in the number of transistors per chip as well as improved energy efficiency per FLOPS. This spurred a dramatic growth in aggregate computational performance of the largest supercomputing systems, yielding multiple Petascale implementations deployed in various locations over the world. Unfortunately, these advances did not translate to the required extent into accompanying I/O systems, which primarily saw the improvement in cumulative storage sizes required to match the ever expanding volume of scientific data sets, but little more in terms of architecture or effective access latency. Moreover, while new models of computations are formulated to handle the burden of efficiently structuring the parallel computations in anticipation of the arrival of Exascale systems, a meager progress is observed in the area of storage subsystems. New classes of algorithms developed for massively parallel applications, that gracefully handle the challenges of asynchrony, heavily multithreaded distributed codes, and message-driven computation, must be matched by similar advances in I/O methods and algorithms to produce a well performing and balanced supercomputing system. This paper discusses PXFS, a file system model for persistent objects inspired by the ParalleX model of execution that addresses many of these challenges. An early implementation of PXFS utilizing a well known Orange parallel file system as its back-end via asynchronous I/O layer is also described along with the preliminary performance data. The results show perfect scalability and 3x to 20x times speedup of I/O throughput performance comparing to OrangeFS user interface. Also the PXFS module on OrangeFS with 24 clients sees a 5x to 10x times more throughput than NFS.

I. INTRODUCTION

Over the coming decade the performance of the largest computing systems is expected to shift from Terascale and Petascale to Exascale. According to the TOP500 supercomputer list [1] of June 2013, the performance of the top 26 systems breaks the PetaFLOPs barrier. These systems are composed of tens of thousands of cores and nodes running in parallel. For example, the top supercomputer, Tianhe-2, has 16,000 nodes equipped with 3,120,000 computing cores and achieves the performance of 33.86 petaFLOP/s on the Linpack benchmark [1]. The rapid growth in computing cores and high demand for parallelism imposes significant challenges for the parallel runtime and storage system.

Applications employing conventional parallel runtime systems, such as the Communicating Sequential Processes (CSP) [2] execution model as reflected by the Message Passing

Interface (MPI) [3] programming model, are getting more and more difficult to make effective use of the ever increasing number of processors to achieve desired scalability and performance. The main limiting factors are: a) *Starvation* (insufficient concurrent work to maintain high utilization of resources), b) *Latencies* (delay of remote resource access and services), c) *Overheads* (work for management of parallel actions and resources on critical path which is not necessary in a sequential variant), and d) *Waiting for contention resolution* (delays due to lack of availability of oversubscribed shared resources). All of these factors (**SLOW**) are difficult to avoid using today's prevalent programming models, and a new computational strategy is required to achieve dramatic increases in performance. The ParalleX execution model [4]–[6] is offered as a means of addressing these critical computational requirements. It is striving to expose myriad forms of parallelism, hide system wide latencies, decouple hardware execution resources from executing software tasks to prevent the blocking of processor cores, and to enable runtime dynamic adaptive scheduling to employ real-time system state to resource management decisions.

In the meantime, parallel storage system is expected to handle the input/output (I/O) requests from parallel applications with good performance and great scalability. Many parallel file systems are developed to answer that challenge [7]–[10]. Orange File System (OrangeFS) [8], [11], [12] is a production-quality parallel file system designed for use on high end computing (HEC) systems that provides very high performance access to disk storage for parallel applications. However, I/O for Exascale high-end computing (HEC) systems is still hampered by several unfortunate issues, including orders of magnitude slower speed and response time, storage distribution problem, fast growing volume of application data and complex data structures.

The challenges of Exascale computing suggest that not a simple extension of our current model of computation, but rather a new model of computation is needed upon which a new framework for mass storage may be built. It is both prudent and essential to consider a corresponding model of I/O along with the design of next generation computation ecosystem. In this paper we present PXFS (ParalleX File System), a new I/O model aiming for Exascale computing.

It will be taking advantage of ParalleX execution model and OrangeFS parallel storage research to extract the maximum of parallelism and performance out of the storage resources. Based on some preliminary work [13], [14], this paper, for the first time, presents a novel design of PXFS model and a complete implementation using HPX and OrangeFS interfaces. This paper also evaluates its performance against traditional model and shows tens of times speedup on I/O throughputs.

In the rest of the paper, the related work is listed in Section II. The design and a primitive implementation of PXFS is detailed in Section III. The performance evaluation is presented in Section IV. Finally Section V draws the conclusion and a plan of future work.

II. BACKGROUND

High Performance ParalleX (HPX) [5], [15]–[17] is the first open-source implementation of the ParalleX execution model. HPX is a state-of-the-art runtime system developed for conventional architectures and, currently, Linux-based systems, such as large Non Uniform Memory Access (NUMA) machines and clusters. The modular framework facilitates simple compile- or runtime-configuration and minimizes the runtime footprint. HPX has been carefully designed as an alternative to mainstream parallel frameworks such as MPI. It focuses on overcoming conventional limitations such as global barriers, poor latency hiding, and lack of support for fine-grained parallelism.

The current implementation of HPX (see Fig. 1) supports most of the key ParalleX elements: Parcels, PX-threads, Local Control Objects (LCOs) and the Active Global Address Space (AGAS).

HPX currently implements AGAS as a set of services that support a 128-bit global address space spanning all localities. AGAS provides two naming layers in HPX. The primary naming service maps 128-bit unique, global identifiers (GIDs) to a tuple of meta-data that can be used to locate an object on a particular locality. The higher-level layer maps hierarchical symbolic names to GIDs. Unlike systems such as X10 [18], Chapel [19], or UPC [20], which are based on PGAS [21], AGAS exposes a dynamic, adaptive address space which evolves over the lifetime of an HPX application. When a globally named object is migrated, the AGAS mapping is updated, however, its GID remains the same. This decouples references to those objects from the locality that they are located on.

LCOs provide a means of controlling parallelization and synchronization in HPX. Any object that may create a new HPX-thread or reactivates a suspended HPX-thread exposes the required functionality of an LCO. Support for event-driven HPX-thread creation, protection of shared data structures, and organization of flow control are provided by LCOs. They are designed to allow for HPX-threads to proceed in their execution as far as possible, without waiting for a particular blocking operation, such as a data dependency or I/O, to finish. Some of the more prominent LCOs provided by HPX are:

- *Futures* [22]–[24] are proxies for results that are not yet known, possibly because they have not yet been computed. A future synchronizes access to the result value associated with it by suspending HPX-threads requesting the value if the value is not available at the time of the request. When the result becomes available, the future reactivates all suspended HPX-threads waiting for the value. These semantics allow execution to proceed unblocked until the actual value is required for computation.
- *Dataflow objects* [25]–[27] provide a powerful mechanism for managing data dependencies without the use of global barriers. A dataflow LCO waits for a set of values to become available and triggers a predefined function passing along all input data.
- *Traditional concurrency control mechanisms* such as various types of mutexes [28], counting semaphores, spinlocks, condition variables and barriers are also exposed as LCOs in HPX. These constructs can be used to cooperatively block an HPX-thread while informing the HPX thread-manager that other HPX-threads can be scheduled on the underlying OS-thread.

The HPX thread-manager is responsible for the creation, scheduling, execution and destruction of HPX-threads. In HPX, threading uses an $M : N$ or hybrid threading model. In this model, N HPX-threads are mapped onto M kernel threads (OS-threads), usually one OS-thread per core. This threading model was chosen to enable fine-grained parallelization and low overhead context switches; HPX-threads can be scheduled without a kernel call, reducing the overhead of their execution and suspension. The thread-manager uses a work-queue based execution strategy with work stealing similar to systems such as Cilk++ [29], Intel Threading Building Blocks (TBB [30]) and the Microsoft Parallel Patterns Library (PPL [31]). HPX-threads are scheduled cooperatively, that is, they are not preempted by the thread-manager. HPX-threads voluntarily suspend themselves when they must wait for data that they require to continue execution, I/O operations, or synchronization. The latter fact requires special consideration when designing a persistent storage model as care must be taken not to suspend any of the HPX-threads on the OS level. This would block any further progress as no other HPX-threads could be executed while the underlying OS thread is suspended.

The software architecture of OrangeFS is shown in Fig. 2. It deploys a client/server structure and dynamically distributes file data and metadata across a system. This strategy alleviates file system bottlenecks and improves system scalability. There is a user library which provides system call APIs, POSIX library APIs to be incorporated to other applications.

III. DESIGN AND IMPLEMENTATION

The PXFS model is designed as a layer on top of storage media and file systems, and as a part of a parallel runtime system. The system diagram is depicted in Fig. 3. It is implemented as a component of HPX and manages I/O operations through different levels.

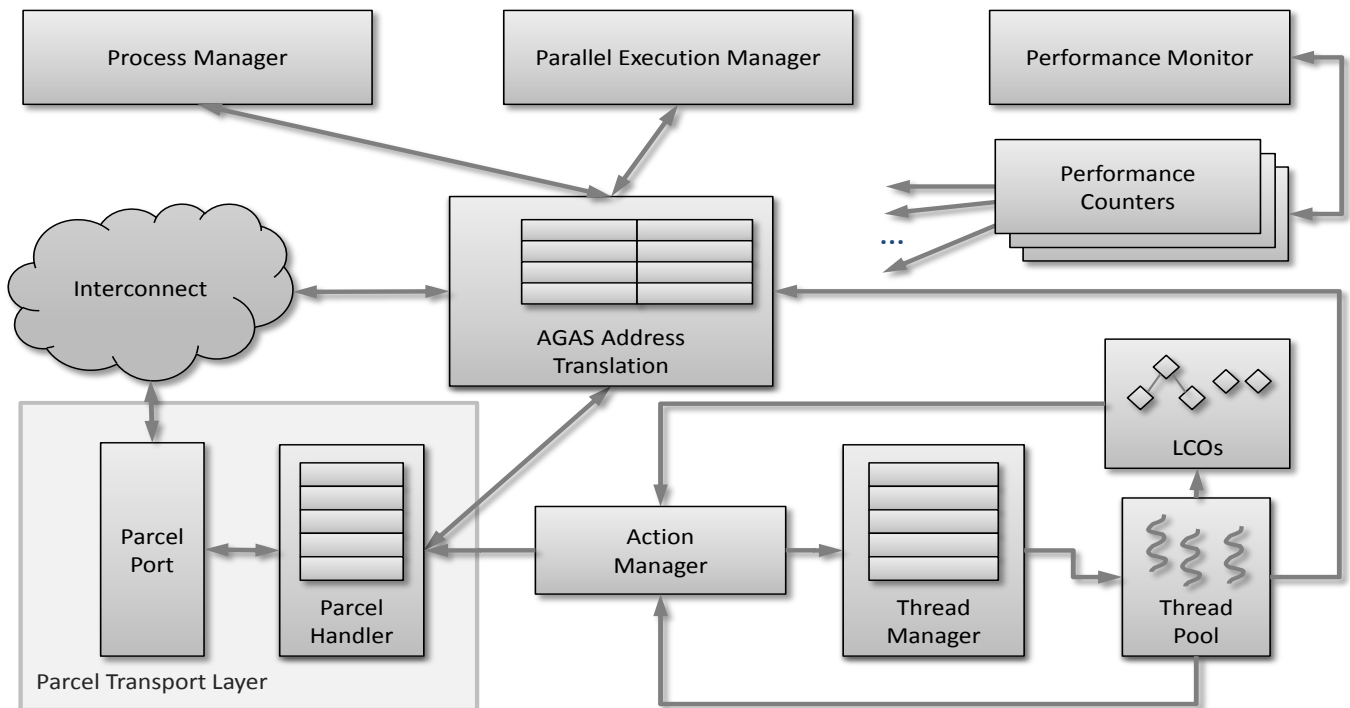


Fig. 1. Architecture of the HPX runtime system. HPX implements the supporting functionality for most of the elements needed for the ParalleX model: Parcels (parcel-port and parcel-handlers), HPX-threads (thread-manager), LCOs, AGAS, HPX-processes, performance counters and a means of integrating application specific components.




Fig. 2. Software architecture of OrangeFS parallel file system.

While designing the PXFS APIs, special attention was directed towards a natural integration with the existing highly asynchronous programming model exposed by HPX. By design, all functionality in HPX which can potentially take longer than 100 microseconds to execute is exposed through asynchronous functions. An asynchronous function in HPX is a function which is returning a future representing the result of that function. The caller can continue executing immediately without being suspended or without having to wait for the function result. The code example in Fig.4 shows part of the developed file oriented API. Note that the HPX future type exposes an interface consistent with the C++11 Standard [32] with extensions as proposed to the C++ standardization process [33], [34].

The advantage of the chosen asynchronous programming model for PXFS is demonstrated in the code example shown

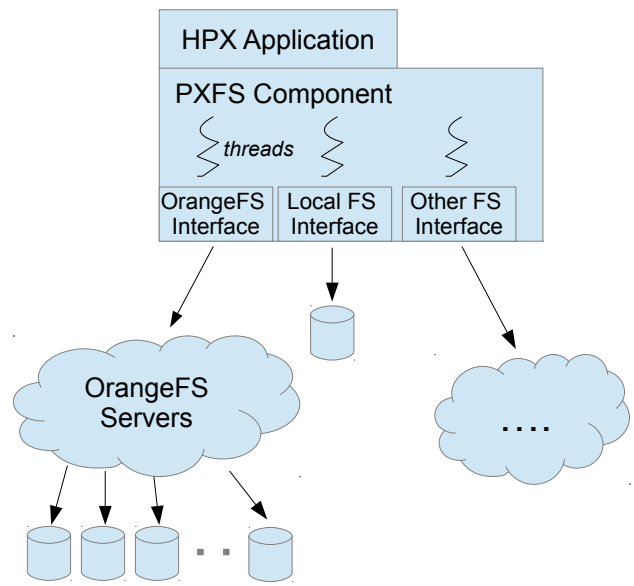


Fig. 3. Design of PXFS storage component as a part of HPX runtime.

in Fig.5, which demonstrates how easy it is to fully overlap the asynchronous I/O operations (create a file followed by writing a chunk of data to it) with other useful work.

As a coordinated work, an asynchronous interface is also designed and a prototype has been implemented [14] in OrangeFS. It is part of the PXFS model to interact with parallel file system in an asynchronous fashion. The asynchronous

```

Exemplar Asynchronous File API
namespace hpx { namespace io {
struct file {
    ~file();           // closes file synchronously

    static future<file> create(std::string name, int mode);
    static future<file> open(std::string name, int mode);

    future<void> close();

    // read/write operations
    future<ssize_t> read(const void *buf, size_t count);
    future<ssize_t> write(const void *buf, size_t count);
    future<off_t> lseek(off_t offset, int whence);
};
}}

```

Fig. 4. This code example shows the asynchronous file oriented API exposed by the implementation of the presented persistent storage model.

```

Exemplar Asynchronous File API
using namespace hpx;
using namespace hpx::io;

std::vector<char> data = {...};

// spawn asynchronous file creation and write operation
future<void> op =
    file::create("filename", mode).then(
        [&data](future<file> f) {
            f.get().write(data);
        });

// do other operations here, concurrently to I/O
op.wait(); // synchronize with whole I/O operation

```

Fig. 5. This code example demonstrates an exemplar use of the file oriented API to asynchronously create a new file and to write some data to it. The I/O operation is performed fully overlapping any other work which needs to be performed before synchronizing with the result of the I/O.

interface will use the OrangeFS system calls to schedule I/O operations and execute a callback function when the task is finished.

The implementation of PXFS will create HPX futures for I/O operations and set the future value in the callback function of OrangeFS asynchronous APIs to integrate the parallel file system into the HPX runtime system seamlessly. Thus the PXFS APIs provide a uniform asynchronous interface to bridge the gap between HPX runtime components and OrangeFS storage components efficiently and transparently. The nature of object oriented programming can enable other file systems to be adopted into PXFS model easily.

IV. PERFORMANCE EVALUATION

A. Method

A disk performance micro benchmark is developed in HPX to measure the total throughput of PXFS module on I/O reading and writing performance. The benchmark can adjust the following parameters:

- nc : number of client nodes
- nt : number of threads on each client node, a thread is carrying out the actual I/O operations.
- r/w : select between read and write tests.
- nf : number of test files.

- nb : number of total blocks composed in one file. A thread is operating on one block when issuing one I/O commands.
- sb : number of bytes in one block.

These test files are read or written in a sequential fashion at the current time. Then the read/write throughput can be calculated as

$$\text{Throughput} = \frac{nc * nt * nf * nb * sb}{\text{Elapsed Time}}$$

It is obvious that larger throughput means a better performance.

B. Experiment Setup

All the tests are executed on the Hermione cluster at STE||AR [16] group in Center for Computation and Technology (CCT) [35] at Louisiana State University (LSU) [36]. At the time when the experiments are conducted, the cluster is a heterogeneous system consisting of 39 computing nodes connected with Gigabit Ethernet. The nodes are running Linux and using SLURM as the scheduler.

Four OrangeFS systems are started with 2,4,8,16 server nodes. The client nodes are distributed in the cluster and not using the server nodes. The number of client nodes is selected to be comparable to the number of server nodes. HPX runtime system is responsible for managing threads on the client nodes. Large number of blocks and files are picked to keep the system under its full capacity.

Several sets of test cases is performed in this experiment on HPX with and without PXFS module systems on top of OrangeFS file system and Network File System (NFS) [37].

- w-5f-100x64KB: writing 5 files with 100 blocks and block size 64KB;
- r-5f-100x64KB: reading 5 files with 100 blocks and block size 64KB;
- w-5f-100x1MB: writing 5 files with 100 blocks and block size 1MB;
- r-5f-100x1MB: reading 5 files with 100 blocks and block size 1MB;

These four file test cases can represent reading and writing medium and large block size cases in large scale scientific applications which read and create sets of data frequently.

As an alternative use case, the OrangeFS user interface is called directly from the benchmark. The user interface can only support 1 thread per client node and deploys a synchronous API set. In the PXFS module, multiple threads can be enabled and multi-core architecture can be utilized.

C. Results and Discussions

The I/O throughput results of four test cases are presented in Fig. 6, Fig. 7, Fig. 8 and Fig. 9 respectively.

It is crystal clear that the PXFS module has increased the throughput dramatically than the direct synchronous interface. For one thread, 24 clients and 16 OrangeFS servers, the PXFS has 10x times speedup for w-5f-100x64KB, 20x times speedup for r-5f-100x64KB, 4x times speedup for w-5f-100x1MB and 3x times speedup for r-5f-100x1MB.




Fig. 6. I/O performance of file test case $w-5f-100 \times 64KB$ between PXFS and direct interface on HPX and OrangeFS.




Fig. 9. I/O performance of file test case $r-5f-100 \times 1MB$ between PXFS and direct interface on HPX and OrangeFS.

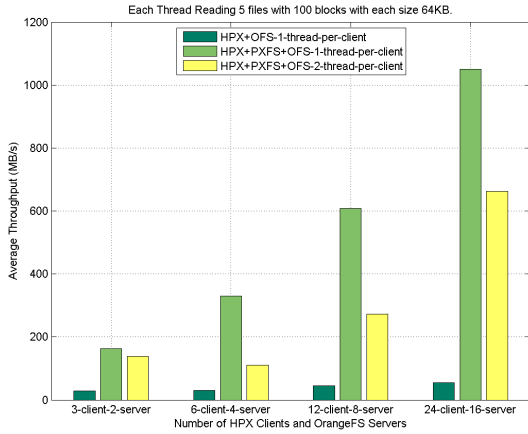


Fig. 7. I/O performance of file test case $r-5f-100 \times 64KB$ between PXFS and direct interface on HPX and OrangeFS.

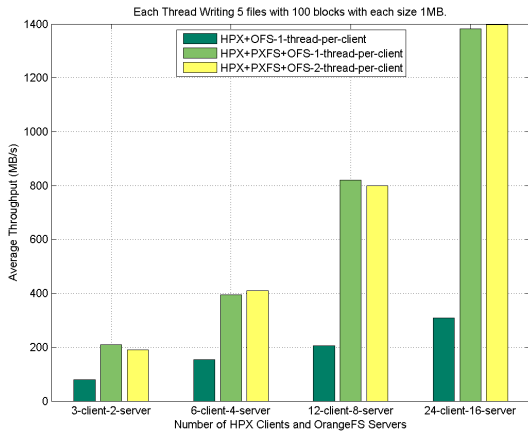


Fig. 8. I/O performance of file test case $w-5f-100 \times 1MB$ between PXFS and direct interface on HPX and OrangeFS.

When using multiple threads, $w-5f-100 \times 64KB$ nearly doubles the throughput, $r-5f-100 \times 64KB$ sees some downgrades and $w/r-5f-100 \times 1MB$ have little variations. It is indicated that the thread management and I/O access pattern can influence the throughput and the $w-5f-100 \times 64KB$ case still has potentials for higher performance, while $r-5f-100 \times 64KB$ might have some thread contention and needs some fine tune on I/O and parallelism parameters. More work will be done to examine the PXFS performances with large number of threads and different I/O access patterns.

In all of the four cases, PXFS has shown perfect scalability when the number of OrangeFS servers is changing from 2 through 16. The throughput increases at the same ratio as the number of OrangeFS servers. The great scalability is a strong argument for running large scale application on thousands of nodes now and in the future.

To compare the performance of a parallel file system and a centralized file system, the tests are also run on NFS with 24 HPX clients and the results are displayed in Fig. 10.

As seen from the results, the performance with the PXFS model and a parallel system underneath greatly surpasses the NFS system with a 5x to 10x times improvement in I/O throughput with 24 HPX clients. It is shown that in Exascale era, the new storage model should take into account the experience of parallel file system development.

V. CONCLUSION

This paper presents PXFS, a novel persistent storage model for extreme scale. It aims to explore the maximum of parallelism and performance in the Exascale computing era, which is expected to arrive in the next decade. The early implementation of PXFS utilizes Orange parallel file system as the back-end and incorporates an asynchronous I/O interface into HPX, an implementation of ParalleX execution model. The I/O performance is evaluated with a homemade benchmark and the I/O throughput shows a perfect scalability, along with a 3x to 20x times speedup against the OrangeFS user interface

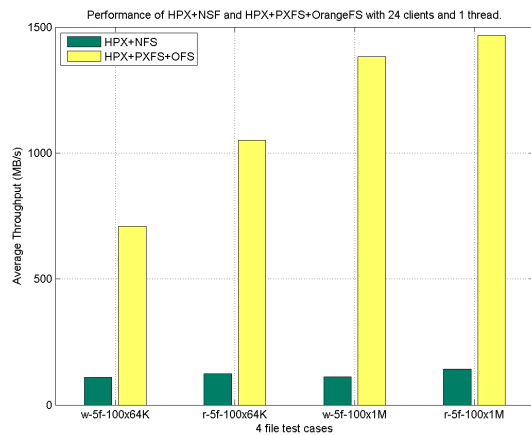


Fig. 10. I/O performance of all four file test cases between HPX+PXFS module on OrangeFS and HPX on NFS.

and a 5x to 10x times higher throughput than NFS with 24 clients.

PXFS is our first try to combine the strength of parallel file system and parallel runtime system to extract the maximum of parallelism and performance out of the storage resources. In the future, more layers of Orange parallel file system and ParalleX model will be analyzed to develop a more sophisticated storage model. For the benchmark, the effect of threads and various architectures on I/O performance will be tested thoroughly. Real world applications will be used as another benchmark to evaluate the performance of PXFS.

ACKNOWLEDGMENT

The authors would like to acknowledge support in part from NSF grant IIS-1252358.

REFERENCES

- [1] "TOP500 Supercomputing Sites," <http://www.top500.org>, November 2010.
- [2] S. D. Brookes, C. A. R. Hoare, and A. W. Roscoe, "A theory of communicating sequential processes," *J. ACM*, vol. 31, no. 3, pp. 560–599, 1984.
- [3] Message Passing Interface Forum, *MPI: A Message-Passing Interface Standard, Version 2.2*. Stuttgart, Germany: High Performance Computing Center Stuttgart (HLRS), September 2009.
- [4] G. Gao, T. Sterling, R. Stevens, M. Hereld, and W. Zhu, "ParalleX: A study of a new parallel computation model," in *Parallel and Distributed Processing Symposium, 2007. IPDPS 2007. IEEE International*, 2007, pp. 1–6.
- [5] H. Kaiser, M. Brodowicz, and T. Sterling, "ParalleX: An advanced parallel execution model for scaling-impaired applications," in *Parallel Processing Workshops*. Los Alamitos, CA, USA: IEEE Computer Society, 2009, pp. 394–401.
- [6] A. Tabbal, M. Anderson, M. Brodowicz, H. Kaiser, and T. Sterling, "Preliminary design examination of the ParalleX system from a software and hardware perspective," *SIGMETRICS Performance Evaluation Review*, vol. 38, p. 4, Mar 2011.
- [7] "Lustre a Network Clustering FS," <http://www.lustre.org>.
- [8] "Orange File System," <http://orangefs.org>.
- [9] F. Schmuck and R. Haskin, "GPFS: A Shared-Disk File System for Large Computing Clusters," in *Proceedings of the Conference on File and Storage Technology (FAST'02)*, Berkeley, CA, USA, January 2002.
- [10] P. H. Carns, W. B. L. III, R. B. Ross, and R. Thakur, "PVFS: A Parallel File System for Linux Clusters," in *Proceedings of the 4th annual Linux Showcase & Conference*. Berkeley, CA, USA: USENIX Association, 2000.
- [11] S. Roberts, "Next generation storage for the hltcoe," 2013.
- [12] S. Yang, W. LIGON, and E. C. Quarles, "Scalable distributed directory implementation on orange file system," *Proc. IEEE Intl. Wrkshp. Storage Network Architecture and Parallel I/Os (SNAPI)*, 2011.
- [13] S. Yang, W. L. III, M. Brodowicz, and H. Kaiser, "PXFS: A Persistent Storage Model for Extreme Scale," in *Scientific Computing Around Louisiana*, February 2013.
- [14] S. Snyder, "ParalleX file system (PXFS): Bridging the Gap Between Exascale Processing Capabilities and I/O Performance," Master's thesis, Clemson University, Clemson, SC, USA, May 2013.
- [15] H. Kaiser, B. Adelstein-Lelbach *et al.*, "HPX source code repository," 2007-2013, available under the Boost Software License (a BSD-style open source license). [Online]. Available: <http://github.com/STELLAR-GROUP/hpx>
- [16] The STE||AR Group, "Systems Technologies, Emerging Parallelism, and Algorithms Research," 2011-2013, <http://stellar.cct.lsu.edu>. [Online]. Available: {<http://stellar.cct.lsu.edu>}
- [17] C. Dekate, M. Anderson, M. Brodowicz, H. Kaiser, B. Adelstein-Lelbach, and T. L. Sterling, "Improving the scalability of parallel N-body applications with an event driven constraint based execution model," *Accepted, The International Journal of High Performance Computing Applications*, vol. abs/1109.5190, 2012, <http://arxiv.org/abs/1109.5190>.
- [18] P. Charles, C. Grothoff, V. Saraswat, C. Donawa, A. Kielstra, K. Ebcioglu, C. von Praun, and V. Sarkar, "X10: An object-oriented approach to non-uniform cluster computing," *SIGPLAN Not.*, vol. 40, pp. 519–538, October 2005. [Online]. Available: <http://doi.acm.org/10.1145/1103845.1094852>
- [19] B. L. Chamberlain, D. Callahan, and H. P. Zima, "Parallel programmability and the Chapel language," *International Journal of High Performance Computing Applications*, vol. 21, pp. 291–312, 2007.
- [20] UPC Consortium, "UPC Language Specifications, v1.2." Lawrence Berkeley National Lab, Tech Report LBNL-59208, 2005. [Online]. Available: <http://www.gwu.edu/~upc/publications/LBNL-59208.pdf>
- [21] PGAS, "PGAS - Partitioned Global Address Space," 2011, <http://www.pgms.org>. [Online]. Available: {<http://www.pgms.org>}
- [22] H. C. Baker and C. Hewitt, "The incremental garbage collection of processes," in *SIGART Bull.* New York, NY, USA: ACM, August 1977, pp. 55–59. [Online]. Available: <http://doi.acm.org/10.1145/872736.806932>
- [23] D. P. Friedman and D. S. Wise, "Cons should not evaluate its arguments," in *ICALP*, 1976, pp. 257–284.
- [24] R. H. Halstead, Jr., "MULTILISP: A language for concurrent symbolic computation," *ACM Trans. Program. Lang. Syst.*, vol. 7, pp. 501–538, October 1985. [Online]. Available: <http://doi.acm.org/10.1145/4472.4478>
- [25] J. B. Dennis, "First version of a data flow procedure language," in *Symposium on Programming*, 1974, pp. 362–376.
- [26] J. B. Dennis and D. Misunas, "A preliminary architecture for a basic data-flow processor," in *25 Years ISCA: Retrospectives and Reprints*, 1998, pp. 125–131.
- [27] Arvind and R. Nikhil, "Executing a program on the MIT tagged-token dataflow architecture," in *PARLE '87, Parallel Architectures and Languages Europe, Volume 2: Parallel Languages*, J. W. de Bakker, A. J. Nijman, and P. C. Treleaven, Eds. Berlin, DE: Springer-Verlag, 1987, lecture Notes in Computer Science 259.
- [28] P. J. Courtois, F. Heymans, and D. L. Parnas, "Concurrent control with "readers" and "writers"," *Commun. ACM*, vol. 14, no. 10, pp. 667–668, 1971.
- [29] C. E. Leiserson, "The Cilk++ concurrency platform," in *DAC '09: Proceedings of the 46th Annual Design Automation Conference*. New York, NY, USA: ACM, 2009, pp. 522–527. [Online]. Available: <http://dx.doi.org/10.1145/1629911.1630048>
- [30] Intel, "Intel Thread Building Blocks 3.0," 2010, <http://www.threadingbuildingblocks.org>. [Online]. Available: <http://www.threadingbuildingblocks.org/>
- [31] Microsoft, "Microsoft Parallel Pattern Library," 2010, <http://msdn.microsoft.com/en-us/library/dd492418.aspx>. [Online]. Available: <http://msdn.microsoft.com/en-us/library/dd492418.aspx>
- [32] The C++ Standards Committee, "ISO/IEC 14882:2011, Standard for Programming Language C++," Tech. Rep., 2011,

std.org/jtc1/sc22/wg21. [Online]. Available: {<http://www.open-std.org/jtc1/sc22/wg21>}

- [33] Niklas Gustafsson and Artur Laksberg and Herb Sutter and Sana Mithani, "Improvements to std::future;T; and Related APIs," Tech. Rep., 2013, <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2013/n3634.pdf>. [Online]. Available: {<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2013/n3634.pdf>}
- [34] Peter Dimov, "Additional std::async Launch Policies," Tech. Rep., 2013, <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2013/n3632.html>. [Online]. Available: {<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2013/n3632.html>}
- [35] "Center for Computation & Technology," <http://cct.lsu.edu>.
- [36] "Louisiana State University," <http://www.lsu.edu>.
- [37] S. Shepler, B. Callaghan, D. Robinson, R. Thurlow, C. Beame, M. Eisler, and D. Noveck, "Network file system (nfs) version 4 protocol," 2003.