



Maximum Performance Computing Solutions



Maxeler makes dataflow supercomputers

Our customers have achieved order-of-magnitude speed-ups

in applications as diverse as seismic modeling for the oil and gas industry, risk calculation for financial institutions and genome sequencing for cutting-edge healthcare. Improvements on this scale have impact beyond speeding up applications: they transform businesses.

What makes these supercomputers so much faster is our Dataflow technology:

a Dataflow supercomputer is fundamentally different because it changes its structure to implement an optimised computer for the current application. Our supercomputers are easily deployed within existing high-



Contents

Welcome.....	2
Dataflow Computing.....	4
Dataflow Engines	5
Software.....	6
MPC-X Series	8
MPC-N Series	10
MPC-C Series.....	12
Desktop	13
MaxCloud	13
Consultancy.....	14
Get In Touch	15

Oil & Gas
Financial Analytics
Low-Latency Networking
Life Sciences
Big Data
Scientific Computing



performance infrastructure, either as stand-alone CPU+Dataflow nodes or as network-connected Dataflow nodes is an existing cluster. Our management software runs on standard operating systems and seamlessly manages the Dataflow resources.

As our technology is groundbreaking, we also offer software development tools, services and training to get your project off the ground and make sure it is a success. Our multi-disciplinary team of mathematicians, scientists and engineers combines domain expertise, general HPC knowledge and experience with our technology, which allows us to work with you and understand your unique requirements.

“With the new Maxeler technology, J.P. Morgan’s trading businesses can now compute orders of magnitude more quickly, making it possible to improve our understanding and control of the profile of our complex trading risk”

Peter Cherasia,
Head of Markets Strategies,
J.P. Morgan Investment Bank



Dataflow Computing

Maxeler systems exploit dataflow computing – a revolutionary way of performing computation, completely different to computing with conventional CPUs (control flow cores).

Computing with a control flow core

In a software application, the program source is transformed into a list of instructions for a particular processor, which is then loaded into the memory attached to the processor.

Data and instructions are read from memory into the processor core, where operations are performed and the results are written back to memory. Modern processors contain many levels

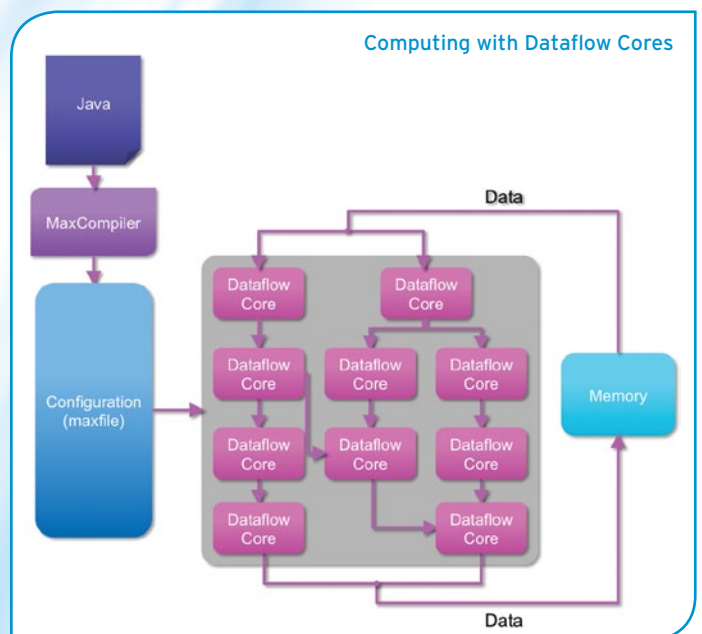
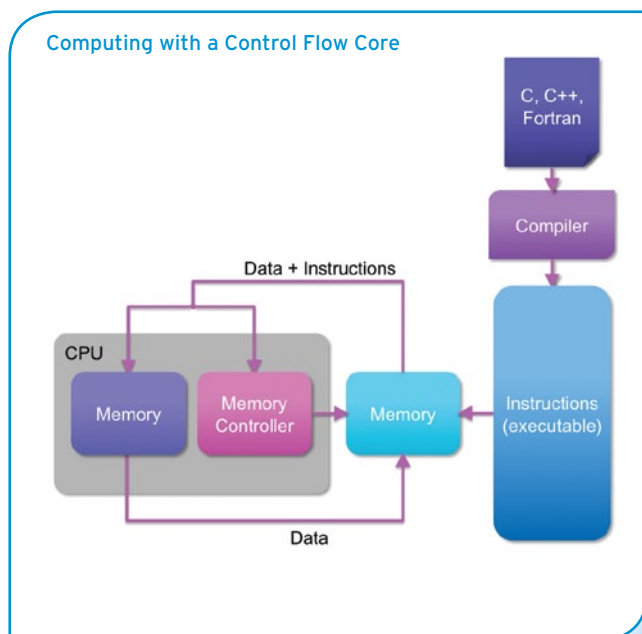
of caching, forwarding and prediction logic to improve the efficiency of this paradigm; however the model is inherently sequential with performance limited by the latency of data movement in this loop.

Computing with dataflow cores


In a dataflow application, the program source is transformed into a dataflow engine configuration file, which describes the operations, layout and connections of a chip specifically

designed for the particular application. Data can be streamed from memory into the chip where operations are performed and data is forwarded directly from one computational unit ('dataflow core') to another, as the results are needed, without being written to the off-chip memory until the chain of processing is complete.


Once a program has finished running, the dataflow engine can be reconfigured for a new application in less than a second.




Dataflow engines enable the computer to be customized to the application



Wave modeling



Monte Carlo pricing




Fluid dynamics




Seismic processing

Speedup for 70Hz FD Seismic Modeling at ENI,
a major energy company



Speedup for Monte Carlo Pricing at J.P. Morgan,
a tier one Investment bank



Dataflow Engines

Maxeler hardware systems include multiple Dataflow Engines (DFEs) providing the dataflow computing resource. Maxeler provides several different DFE architectures for our MPC series and desktop series hardware platforms.

Vectis

- 2,000 parallel multipliers
- 24GB DRAM

Coria

- 3,000 parallel multipliers
- 48GB DRAM

Maia (available Q4 2012)

- 4,000 parallel multipliers
- 48GB DRAM

Software

Maxeler software tools enable scientists and software engineers to rapidly develop and deploy accelerated applications using our high-performance dataflow systems.

MaxCompiler

MaxCompiler is a general-purpose, high-level programming tool suite that provides an intuitive and powerful point of entry for implementation of a wide range of applications as dataflow engines (DFEs).

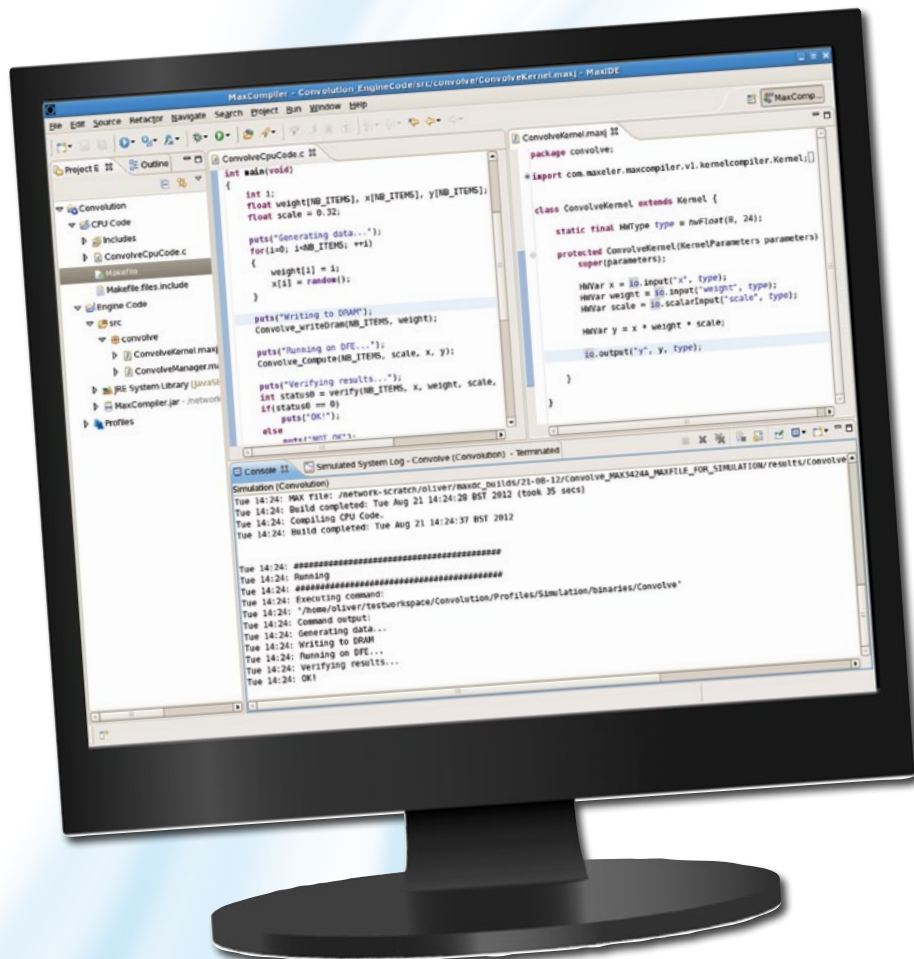
A DFE configuration can be integrated into a CPU executable for deployment on a Dataflow-enabled node or cluster using a simple function API and then linking to a library containing the DFE. This allows DFE configurations to be developed once and reused in a number of applications or distributed for integration into end-user applications. To create applications exploiting new DFE configurations, the MaxCompiler tool-flow splits development into three components:


- Kernel(s), which implement the computational components of the application in hardware.
- Manager configuration, which connects Kernels to the CPU, DFE RAM, other Kernels and other DFEs via the MaxRing interconnect.
- Software, where the DFE is integrated into the original application using a simple API for transferring data to and from the DFE directly and into and out of the RAM of the DFE.

Programmers develop Kernels by writing programs in Java, however only minimal familiarity with Java is required. MaxCompiler includes a complete development environment, MaxIDE, which is based on Eclipse and automates the project management and build processes, greatly accelerating the learning and development process.

During development, the high-speed simulator enables verification of code correctness before generating a hardware implementation.

Once a DFE configuration has been loaded into hardware, the MaxDebug tool allows users to examine the live state of the engine as it runs.






MaxCompilerMPT

MaxCompilerMPT is a complete development environment offering fast and predictable delivery of low-latency networking applications.

Built on top of MaxCompiler, programmers are able to build rich trading applications with built-in exchange connectivity and microsecond-level latency.

Performance and latency can be analyzed with latency annotation - allowing line-by-line latency optimization and management. MaxCompilerMPT incorporates a wide variety of networking features including full UDP and TCP/IP connectivity and can leverage Maxeler's fully-maintained exchange interfaces for rapid development of low latency trading applications.



MaxGenFD

Maxeler MaxGen systems are domain-specific compilers that enable programmers to easily harness the full power of dataflow computing without needing a detailed understanding of the underlying hardware.

MaxGenFD, for 3D finite difference applications, enables geoscientists to engage in rapid development of accelerated seismic processing applications such as Forward Modeling, Reverse Time Migration and Waveform Inversion.

MaxGenFD handles the complexities facing any finite difference implementation such as managing very large data sets, boundary conditions and domain decomposition. In addition, the compiler removes the need for the geoscience programmer to

perform low-level optimizations such as customizing data-types and generating optimized stencil descriptions for hardware.

From a user's high-level FDKernel description, MaxGenFD automatically generates a parallelized implementation, both within a single DFE and across multiple DFEs with halo exchange via the dedicated MaxRing interconnect.

MaxGenFD is implemented as a layer on top of MaxCompiler, the general-purpose dataflow programming system. This allows the user to retain the full programming power of MaxCompiler, while MaxGenFD provides pre-defined libraries to implement common features of finite difference applications.



MPC-X Series

MPC-X Series dataflow nodes provide multiple dataflow engines as shared resources on the network, allowing them to be used by applications running anywhere in a cluster.



Dataflow as a shared resource brings the efficiency benefits of virtualization to dataflow computing, maximizing utilization in a multi-user and multi-application environment such as a private or public cloud service or HPC cluster.

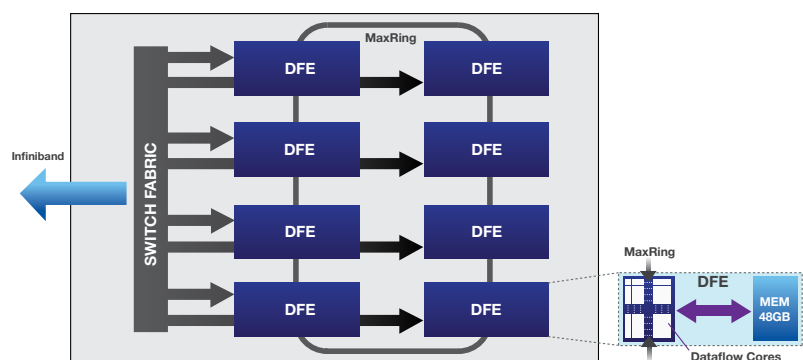
CPU nodes can utilize as many DFEs as are required for a particular application and release DFEs for use by other nodes when not running computation, ensuring all cluster resources are optimally balanced at runtime.

Individual MPC-X nodes provide large memory capacities (up to 384GB) and compute performance equivalent to dozens of conventional x86 servers. The MPC-X series enables remote

access to DFEs by providing dual FDR/QDR Infiniband connectivity combined with Maxeler RDMA technology that supports direct transfers from CPU node memory to remote dataflow engines without inefficient memory copies.

Client CPU nodes can run standard Linux with Maxeler software installed. Maxeler's software automatically manages the resources within the MPC-X nodes, including dynamic allocation of dataflow engines to CPU threads and processes and balancing

MPC-X Series Architecture



demands on the cluster at runtime to maximize overall performance.

With a simple Linux RPM installation, any CPU server in the cluster can be quickly upgraded to begin benefiting from dataflow computing.

An MPC-X node provides eight dataflow engines in a high density 1U form factor with power consumption comparable to a single high-end server.

Each dataflow engine is accessible by any CPU client machine via the Infiniband network, while multiple engines within the same MPC-X node can also communicate directly using the dedicated high-speed MaxRing interconnect.

“An MPC-X node provides eight dataflow engines in a high density 1U form factor with power consumption comparable to a single high-end server.”



Series Features

- Infiniband-connected dataflow engines
- 1U form factor for maximum rack density
- Dual QDR/FDR Infiniband connectivity
- Redundant power supplies
- IPMI/Lights-out management support
- Fully programmable with MaxCompiler
- MaxRing interconnect between DFEs

MPC-X1000

- 8x *vectis* dataflow engines
- Up to 192GB of DFE RAM (24GB per DFE)

MPC-X1200

- 8x *coria* dataflow engines
- Up to 384GB of DFE RAM (48GB per DFE)

MPC-X2000

- 8x *maia* dataflow engines
- Up to 384GB of DFE RAM (48GB per DFE)

MPC-N Series

The MPC-N Series allows ultra low latency line-rate processing of multiple 10Gbit data streams.



An MPC-N Series node provides high-speed network connectivity direct to dataflow engines.

Each node has up to two dataflow engines with a total of 4 SFP/SFP+ ports and 2 CX4 ports, totalling up to 6 high speed connections. Each

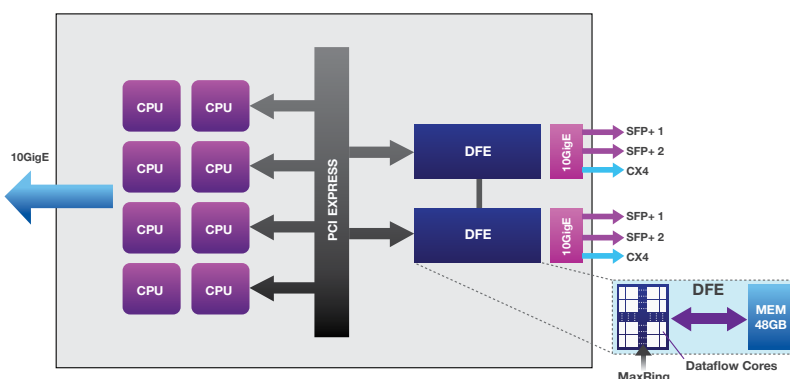
DFE also provides additional inputs for operations such as precision time synchronization from sources including GPS or atomic clocks.

Maxeler DFEs provide the lowest possible latencies while supporting full line-rate processing, with single digit microsecond turnarounds for complex

operations. The MaxCompilerMPT programming environment provides full support for network interfaces, including ultra low latency protocol stacks for common network protocols. Programmers can easily implement the entire latency-sensitive critical path directly in optimized hardware, or can communicate data with legacy applications via low-latency PCI Express data transfer to the CPUs.

The MPC-N Series runs production-standard Linux distributions, including Red Hat Enterprise Linux 5 and 6. The DFE management software installs as a standard RPM package that includes all software necessary to deploy dataflow networking applications, including device driver, monitoring daemon and runtime libraries.

MPC-N Series Architecture





Series Features

- Integrated CPUs and DFEs
- DFEs support multiple 10Gbit data connections
- IPMI/Lights-out management support
- Fully programmable with MaxCompiler

MPC-N40

- **Maximum density connectivity platform**
- 1U form factor
- Up to 2x *vectis* dataflow engines (24GB RAM)
- 2x SFP+ and 1x CX4 connections to each DFE
- 12 Intel Xeon 5600 series CPU cores
- Up to 192GB CPU memory
- Up to 2x 10Gbit Ethernet connections to CPU
- DFEs connect to CPUs via PCI Express gen2 x8
- 5x 2.5" disk drives

MPC-N42

- **Expandable connectivity platform**
- 2U form factor
- Up to 2x *vectis* dataflow engines (24GB RAM)
- 2x SFP+ and 1x CX4 connections to each DFE
- 12 or 16 Intel Xeon E5-2000 series CPU cores
- Up to 384GB CPU memory
- Up to 6x 10Gbit Ethernet connections to CPU
- DFEs connect to CPUs via PCI Express gen2 x8
- 16x 2.5" disk drives
- Redundant power supplies

Financial Exchange Support

Maxeler's fully-maintained Market Data and Order Management libraries offer seamless connectivity to financial venues, allowing developers to focus on developing trading strategies and the core business logic for maximum competitive advantage. Maxeler works with exchanges to certify and maintain exchange connectivity and ensure constant compatibility across major and minor protocol updates. Market Feed Handlers provide multi-level order books and configurable output options, including GPS timestamp, feed information and update fields. Exchange Order Connectivity provides exchange-specific session and order management supporting in-hardware and software order placement, replace and cancel.

“Programmers can easily implement the entire latency-sensitive critical path of an application directly in optimized hardware”

MPC-C Series

MPC-C Series nodes provide the closest coupling of DFEs and CPUs for computations requiring the most data transfer.



An MPC-C Series machine is a 1U server-class HPC system with CPU cores tightly coupled to multiple dataflow compute engines.

Each dataflow engine is connected to the CPUs via PCI Express, and DFEs within the same node are directly connected with MaxRing interconnect.

The node also supports optional Infiniband or 10GE interconnect for tight cluster-level integration.

The MPC-C Series runs standard Enterprise Linux distributions, including Red Hat Enterprise Linux 5 and 6, and are compatible with standard cluster management and job distribution software.

The dataflow engine management software installs as an RPM, and allows applications running on the node to utilize both the dedicated local DFEs and other DFEs shared from MPC-X Series systems elsewhere on the network.



Series Features

- Integrated multi-core CPUs and DFEs
- DFEs connect to CPUs via PCI Express gen2 x8
- 1U form factor
- IPMI/Lights-out management support

MPC-C500

- 4x *vectis* dataflow engines
- Up to 96GB DFE RAM (24GB per DFE)
- 12 Intel Xeon 5600 series CPU cores
- Up to 192GB CPU memory
- MaxRing interconnect between DFEs
- 5x 2.5" disk drives

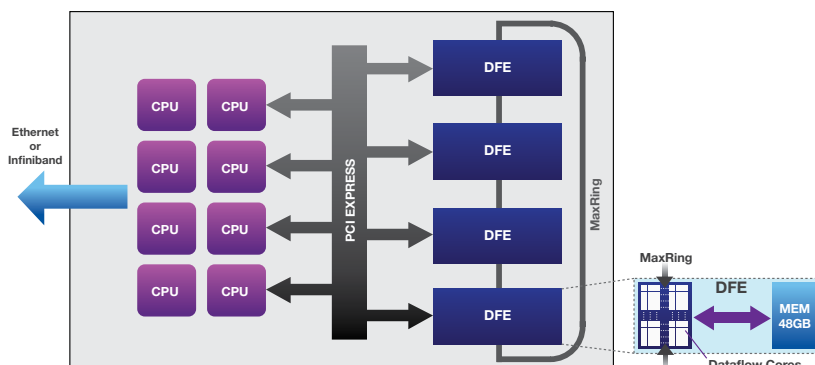
MPC-C600

- 4x *coria* dataflow engines
- Up to 192GB DFE RAM (48GB per DFE)
- 12 Intel Xeon 5600 series CPU cores
- Up to 192GB CPU memory
- 5x 2.5" disk drives

MPC-C755

- 3x *maia* dataflow engines
- Up to 144GB DFE RAM (48GB per DFE)
- 16 Intel Xeon E5-2000 series CPU cores
- Up to 128GB CPU memory
- 4x 2.5" disk drives

MPC-C Series Architecture



Desktop

The MaxWorkstation is a small form factor PC that brings the power of dataflow computing to the desktop.

Small enough to fit on a desk, the MaxWorkstation provides one or two dataflow engines and 4 Intel CPU cores, with a combined processing power equivalent to dozens or hundreds of standard CPU cores.

HPC applications developed on MaxWorkstations are easily deployed into production on MPC-C or MPC-X series data center systems or to MaxCloud, making the MaxWorkstation ideal for software developers as well as users who desire a dedicated compute resource on their desk.

For networking application development, the MaxWorkstation10G provides a desktop development platform with high-speed network connectivity direct to a dataflow engine supporting multiple 10Gbit network connections. Software developed on the MaxWorkstation10G is easily deployed in production on MPC-N series data center connectivity systems.



Series Features

- Small form factor desktop
- Intel Core i7 quad-core CPU with 16GB RAM
- Software portability with data center systems

MaxWorkstation

- **High performance desktop**
- 1 or 2 dataflow engines
- Up to 48GB of RAM per DFE
- DFEs connect to CPU via PCI Express gen2 x8

MaxWorkstation10G

- **Connectivity development platform**
- 1 *vectis* dataflow engine
- 2x SFP+ and 1x CX4 10Gbit connections to DFE
- 2x 10Gbit Ethernet connections direct to CPU
- DFE connects to CPU via PCI Express gen2 x8

MaxCloud

MaxCloud is the industry's first cloud implementation of a high performance dataflow computing system.

MaxCloud offers businesses a scalable off-site resource to eliminate upfront hardware costs and deployment time while benefiting from the high performance of dataflow computing.

MaxCloud compute nodes combine multi-core x86 CPUs with multiple

Maxeler dataflow engines, large memory systems, fast disks and industry-standard Linux.

MaxCloud provides benefits to both new and existing users of Maxeler technology. For those already using Maxeler platforms to run applications, MaxCloud provides a way to provision computation resources only when they are required, and without resorting to capital expenditure. For those new to dataflow computing, MaxCloud allows

the technology to be evaluated, and application development can take place while a hardware purchase is being specified, built and delivered.



Maximum Performance Computing Consultancy


At Maxeler, we use the term Maximum Performance Computing to describe our approach to high-performance dataflow applications.

MPC is a vertical approach to computing, which cuts through the standard layers of abstraction to provide an optimized hardware and software stack for the performance-critical areas of an application.

Maxeler provides MPC consulting and design services to our clients to ensure that projects are efficiently implemented with minimum risk. As a Maxeler client, you will have a specific lead technical and design consultant, who can leverage the multi-disciplinary expertise of the Maxeler team to ensure that your project is a success, managing the Maximum Performance Computing process from start to finish.



Projects often involve accelerating complete applications, beginning with analysis of the problem and any existing solution through to architecting and implementing an optimized solution for the whole application.



Maximum Performance Computing Process

Case Studies

Case Study 1: Oil company

Project duration: 3 months to develop prototype, 3 months productization

Description: Maxeler team developed complete accelerated version for computationally intensive seismic processing application. Finished application was deployed on client's dataflow cluster.

Case Study 2: Investment Bank

Project duration: 9 months

Description: Maxeler team developed dataflow implementations of computational kernels for pricing applications, and worked with client's own development team to integrate into nightly processing flow.

Case Study 3: Oil Services Company

Project duration: 4 months

Description: Client team implemented full application with support from Maxeler

Training

Maxeler runs regular one or two day training courses in dataflow programming at Maxeler offices and partner locations around the world (see our website for upcoming dates). Maxeler can also deliver courses tailored to your specific needs ranging from half day to a full week, either at a Maxeler location or at your site.

Get In Touch

If you'd like to discuss how Maxeler Technologies can help you with your computing challenges, please contact us either by email at info@maxeler.com or by telephone:

US Office


Maxeler Technologies Inc

530 Lytton Ave., Suite 267
Palo Alto,
CA 94301, USA.
Tel: +1 (650) 617 3288

UK Office

Maxeler Technologies Ltd

1 Down Place
London
W6 9JH, UK.
Tel: +44 (0) 208 762 6196



To find out more about our products, services and solutions, please visit www.maxeler.com

Or follow us on twitter  [@MaxelerTech](https://twitter.com/MaxelerTech)

