Symmetry insights for design of supercomputer network topologies: roots and weights lattices.

Alexandre Ferreira Ramos, EACH, USP

School of Arts, Sciences and Humanities - University of São Paulo

May 26, 2014

Applications require further computational capability

- Inertial confinement fusion energy;
- Nuclear fusion and fission science;
- Climate research;
- Design, control and manufacturing of new materials;
- Human brain understanding;

Besides the human inteligence, exascale computing.

Exascale computing – some challenges

At the current typical processing frequency, exascale computing would require $\sim$ billions of processing nodes.

- Power consumption reduction: at current technology an exascale machine would require one Angra 2 facility;
- Dealing with run-time errors: execution and communication;
- Systematic approach to massive parallelism: how to adapt current tools for unprecedent levels of concurrency.

Challenges to enhance communication capability

- Increase communication capabilities to achieve higher performance factor P:

$$P = \frac{used \;\; processing \;\; speed}{nominal \;\; processing \;\; speed}$$

- Hardware level strategy: increase the packing of processing nodes through higher connectivity.

Why increasing the interconnect

- One 24 ports InfiniBand switch power consumption is $\sim$ 300 W.
- One accelerator card power consumption is $\sim$ 300 W.
- Increase of communication has lower power consumption impact than increasing the number of processing nodes.
- Increase of communication might contribute to latency time reduction and, therefore, increasing the P factor of a cluster.

Symmetry

- Symmetry: it is a transformation that preserves invariant a given characteristic of the transformed system.
- Group theory: mathematical framework for quantifying symmetries
- Applications: applications in Physics; Chemistry; Biology; HPC.
- Prototype: rotations on a plan.

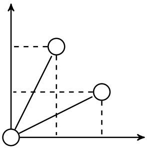

The framework envelops the Theory of Lie Groups; The Lie Algebras; Representation Theory.

Lie algebra – Cartan basis

A denotes a vector space composed by operators A_i with $i = 1, \ldots, n$ over the complex field $\mathbb{C}$. It is a Lie algebra over the complex field if, together with a bilinear map, the Lie bracket

$$A \times A \to A, \quad (A_i, A_j) \mapsto [A_i, A_j] \tag{1}$$

satisfies

$$[A_i, A_i] = 0, \tag{2}$$
$$\left[A_i, \left[A_j, A_k\right]\right] + \left[A_j, [A_k, A_i]\right] + \left[A_k, \left[A_i, A_j\right]\right] = 0, \tag{3}$$

$\forall i, j, k \in \{1, \ldots, n\}$. The square bracket [,] is also denoted the commutator.

Cartan basis and the root vectors

A n-dimensional complex semi-simple Lie algebra can be written in the Cartan basis satisfying the commutation relations:

$$\begin{array}{ll} \left[H_i, H_j\right] = 0 \;\; i,j = 1,\ldots,r, & [E_\alpha, E_{-\alpha}] = \sum \alpha^i H_i \\ [H_i, E_\alpha] = \alpha_i E_\alpha, & [E_\alpha, E_\beta] = N_{\alpha\beta} E_{\alpha+\beta} \end{array} \tag{4}$$

where α indicates a r-dim vector of components α_i, named root vectors. Note that there are $n - r$ root vectors which can be drawn in a diagram named root diagram. r denotes the rank of the algebra and defines the dimension of the root vectors.
Example: $\mathfrak{su}(2)$ algebra (spin) has three operators L_z, L_+, and L_- satisfying:

$$[L_z, L_\pm] = \pm L_\pm, \;\; [L_+, L_-] = 2L_z, \tag{5}$$

Weight diagram

A Lie algebra acts onto a vector space named weight space. The weight vectors Λ's are eigenvectors of the H_i operators (Cartan operators),

$$H_i \left|\Lambda\right\rangle = \Lambda_i \left|\Lambda\right\rangle \tag{6}$$

where Λ_i are the components of the weight vector Λ on a r-dimensional space.

The $\mathfrak{su}(2)$ looks like:

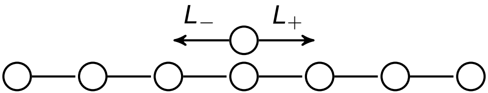

Roots and Weights lattices

A network is designed using a the roots for generating the interconnect topology and the weight vectors generate the nodes.

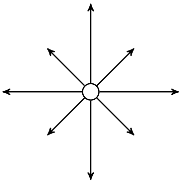

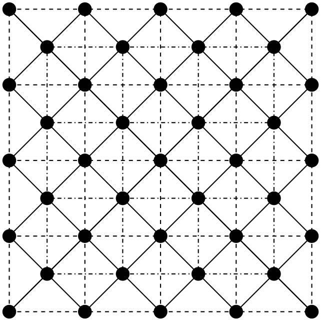

The figures above are representing the network topology generated from the $\mathfrak{sp}(4)$ roots and weights vectors.

Cartan classification – a whole set of network topologies

The Lie algebras are all classified in four different classical families plus five exceptional. The corresponding root diagram for the classical Lie algebras, up to rank 3:

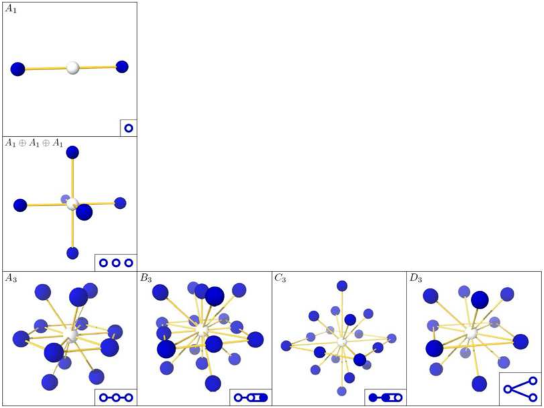

A unifying approach - the blue gene

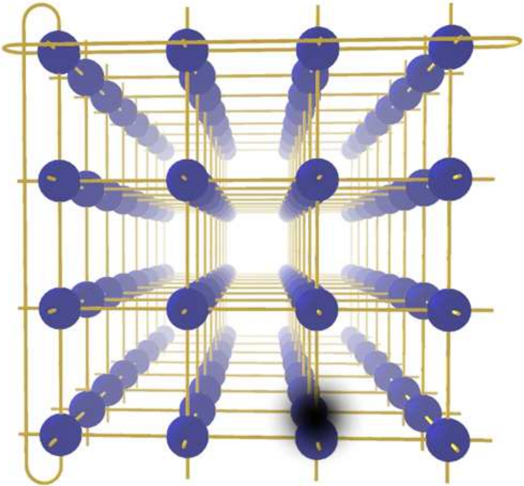

Symmetry breaking – roots

A symmetry breaking chain of the $\mathfrak{sp}(4)$ algebra is:

$$\mathfrak{sp}(4) \supset \mathfrak{sp}(2) \oplus \mathfrak{sp}(2),$$

which is represented in the root diagram as

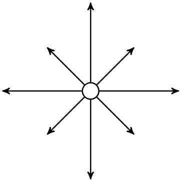

Figure: $\mathfrak{sp}(4)$

$\Longrightarrow$

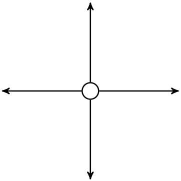

Figure: $\mathfrak{sp}(2) \oplus \mathfrak{sp}(2)$

Symmetry breaking – weights

The $\mathfrak{sp}(4)$ root and weight network will be break into to independent mesh topologies:

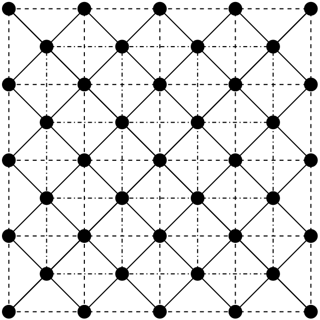

$\Longrightarrow$

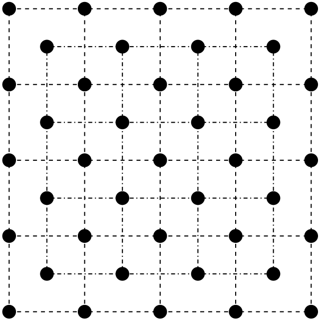

Figure: $\mathfrak{sp}(4)$

Figure: $\mathfrak{sp}(2) \oplus \mathfrak{sp}(2)$

The change could be performed dynamically. The $\mathfrak{sp}(6)$ gives the possibility of $\mu + 1$ $\mathfrak{sp}(4)$ and four 3 dimensional meshes.

Network analysis

- A comparison between hypercubic, mesh, torus and symplectic topologies is presented;
- We take n-dimensional lattices with symplectic $\mathfrak{sp}(2n)$ interconnect topology;
- Number of nodes is indicated by ν and maximum number of edges per node by ϵ;
- Nodes are distributed at lattice built on integers μ^n, $\mu \in \{0, 1, \ldots, \}$;
- The symplectic also have nodes at semi-integers;
- Diameter (L): the maximum distance of the network.

	ν	ϵ	L
Hypercube	2^n	$2n$	n
Mesh	μ^n	$2n$	$n\,\mu$
Torus	μ^n	$2n$	$\lfloor \mu/2 \rfloor$
Symplectic ($\mathfrak{sp}(2n)$)	$(\mu+1)^n + n\,(\mu+1)\mu^{n-1}$	$2n^2$	$n\,\mu$

Hypercube topology

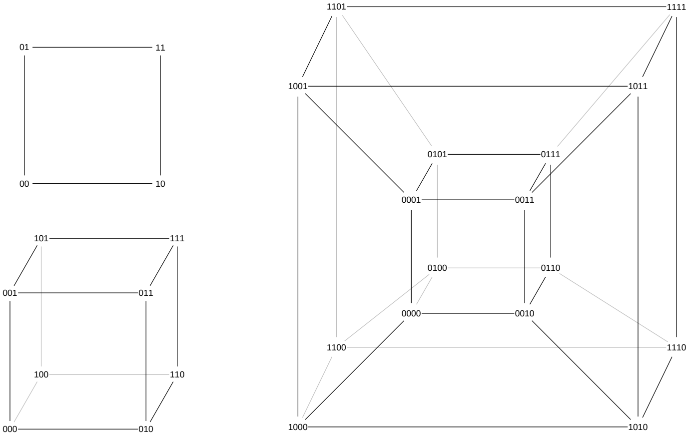

Figure: Hypercubes of, respectively, two, three, and four dimensions.

Mesh topology

A comparison between hypercubic, mesh, torus and symplectic topologies.

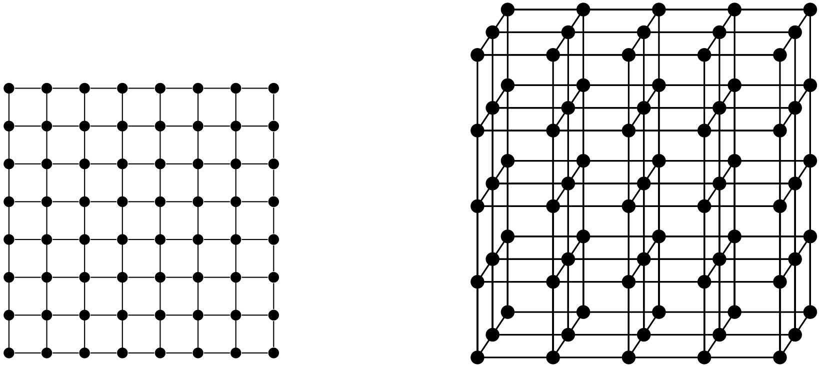

Figure: The two and three dimensional mesh.

Torus topology

A comparison between hypercubic, mesh, torus and symplectic topologies.

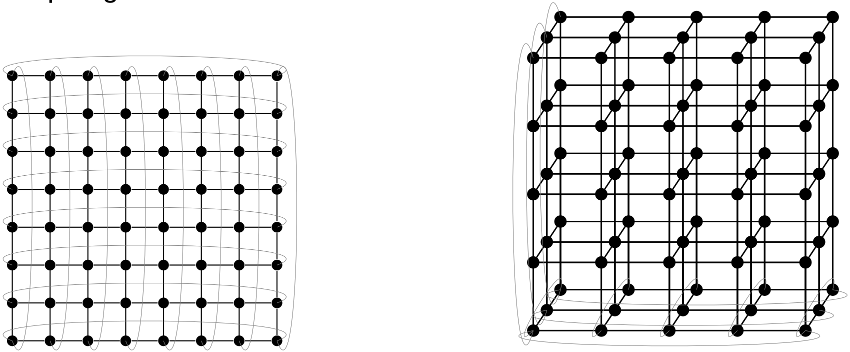

Figure: The two and three dimensional torus.

Distance matrices - Hypercube and Mesh

The matrix element $\Delta_{(i+1),(j+1)}$ indicates the distance between the nodes i and j. The hypercube has distance matrix given by:

$$\Delta_h(n+1) = \begin{bmatrix} \Delta_h(n) & \Delta_h(n) + E(2,n) \\ \Delta_h(2,n) + E^n & \Delta_h(n) \end{bmatrix}, \quad n = 0, 1, \ldots \tag{7}$$

with $\Delta_h(0) = [0]$ and $E(2,n)$ is a $2^n \times 2^n$ matrix with all entries equal to one.

The mesh μ^n has distance matrix $\Delta_m(\mu, n)$:

$$\Delta_m(\mu, n+1) = \begin{pmatrix} \delta_0 & \delta_1 & \ldots & \delta_{\mu-1} \\ \delta_1 & \delta_0 & \ldots & \delta_{\mu-2} \\ \vdots & \vdots & \ddots & \vdots \\ \delta_{\mu-1} & \delta_{\mu-2} & \ldots & \delta_0 \end{pmatrix}, \tag{8}$$

with $\delta_j = \Delta_m(\mu, n) + jE(\mu, n)$.

Distance matrices - Torus and symplectic

The torus has distance matrix

$$\Delta_t(\mu, n+1) = \begin{bmatrix} A(\mu, n+1) & B(\mu, n+1) \\ B^T(\mu, n+1) & A(\mu', n+1,) \end{bmatrix}, \quad n = 0, 1, \ldots \tag{9}$$

with $A(\mu, n+1)$ and $B(\mu, n+1)$ a matrices depending on $\Delta_t(\mu, n)$ and $E(\mu, n)$.

The symplectic distance matrix can be written in terms of the coordinates of the nodes in the lattice

$$\Delta_s(\mu, n+1) = \begin{bmatrix} A_0 & B_1 & \ldots & B_\mu & A_\mu \\ B_1^T & A_0' & \ldots & A_\mu' & B_\mu^T \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ A_\mu & B_\mu^T & \ldots & B_1 & A_0 \end{bmatrix} \tag{10}$$

Distances frequencies for some network topologies

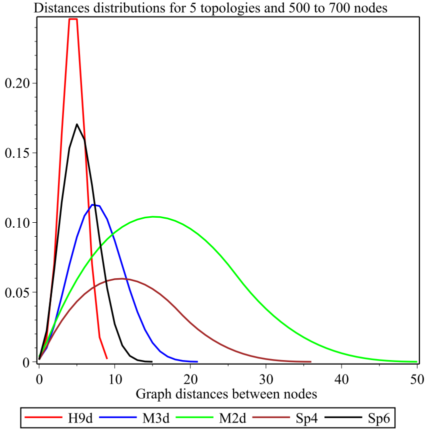

Figure: We compare the histograms for the following topologies: 9-d Hypercube ($\nu = 512$); 2-d, 3-d Mesh($\mu = 26, 8$, $\nu = 676, 512$); $\mathfrak{sp}(4), \mathfrak{sp}(6)$ ($\mu = 18, 5$, $\nu = 685, 666$)

Acknowledgments

Apuama Team members
Prof. Fábio Nakano EACH – USP
Henrique Leme – USP;
Luis Manrique – USP;
Eiji Kawahira – USP;
Carlos Aguni – USP;
Flávio Briz – USP;
Gustavo Burin – USP.

Prof. José Eduardo M. Hornos[1] IFSC – USP;

Prof. Yuefan Deng AMS – Stony Brook University;

[1] Deceased.