Neutral Territory Methods for the Parallel Evaluation of Pairwise Particle Interactions

Kevin J Bowers, Ron O Dror and David E Shaw

D.E. Shaw Research and Development LLC, 39th Floor, Tower 45, 120 West Forty-Fifth Street, New York, NY 10036

Abstract

Particle simulations in fields ranging from biochemistry to astrophysics require evaluation of the interactions between all pairs of particles separated by less than some interaction radius. Recently, Snir [1] and Shaw [2] independently introduced two distinct methods for parallelizing this computation. Both methods achieve asymptotic and practical advantages over traditional parallelization techniques. We describe Shaw's Neutral Territory and Snir's Hybrid Method and show that they represent special cases of a more general class of methods. We also describe new methods that can confer advantages over any previously described method in terms of communications bandwidth and latency. These generalizations include novel spatial decompositions, more advanced import region rounding and multiple zone communication scheduling. Practically speaking, the best choice among the broad category of methods we describe depends on parameters including the interaction radius, the size of the simulated system and the number of processors available. We analyze the best choice of parallelization scheme for different values of these parameters, assuming a simple network cost model. While we do not prove optimality, we show that the communication bandwidth required by several of the schemes comes close to or achieves an approximate lower bound.

[1] Snir, M. *Theory of Computing Systems.* 37, 295-318. 2004
[2] Shaw, D. *Journal of Computational Chemistry.* Accepted for publication. 2005

Cutoff Pairwise Interactions

Rapidly changing short range interactions between close pairs of particles (near electrostatics & van der Waals for example) are often treated as cutoff pairwise interactions.

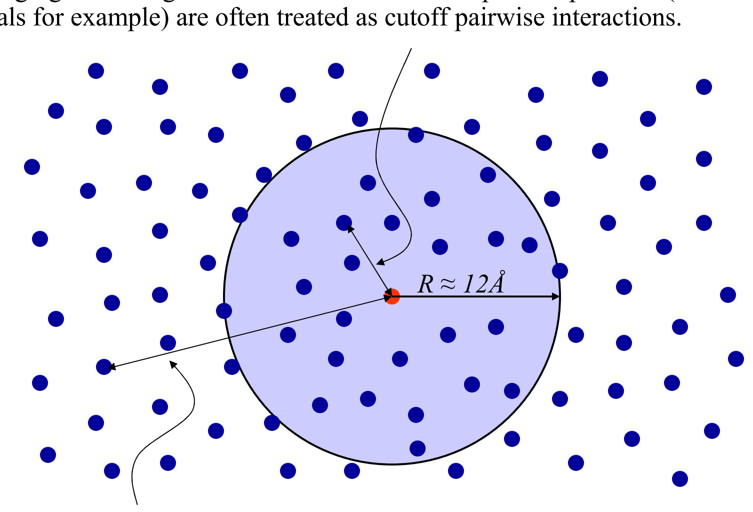

Long-range interactions can be handled by other techniques. See for example the Gaussian Split Ewald method (Shan, Y. Klepeis, J. Eastwood, M. Dror, R. Shaw, D. *Journal of Chemical Physics*. 122, 054101, 2005).

Spatial Domain Decomposition

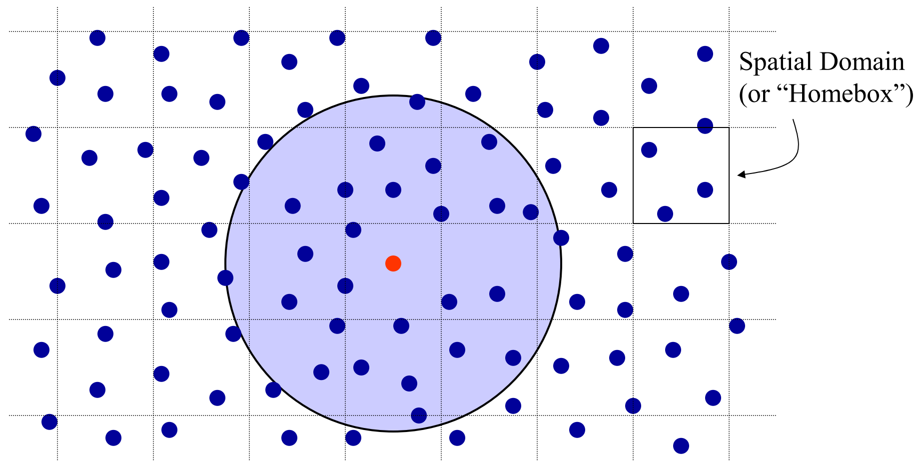

It is often convenient to assign each processor a specific spatial domain (or "homebox"). A homebox is responsible for updating the position of all particles it contains. To compute all cutoff pairwise interactions, particle data will need to be communicated between homeboxes. For fixed simulation volume and particle density, the communication is the limiting factor for large numbers of homeboxes.

Two Strategies for Parallelization

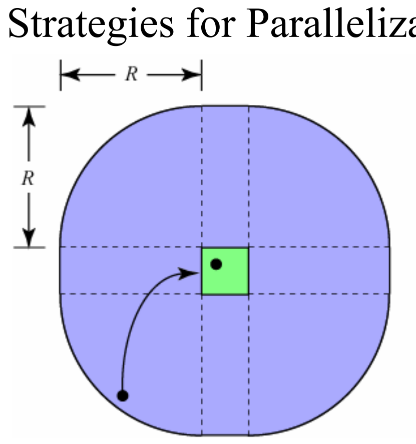

In a traditional method, the interaction is assigned to the homebox that owns at least one of the particles. If the interaction is "commutative"---the interaction (*b-a*)---each homebox will import a "half shell" volume of particle data and a "full shell" if not.

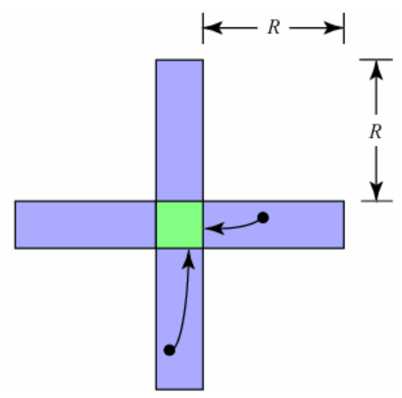

In a "Neutral Territory" method, the interaction may be assigned to a homebox that contains neither particle. That homebox will have to import both particles. The shape of the import region is generally neither a half shell nor a full shell.

Counter-intuitively, NT Methods Improve Bandwidth

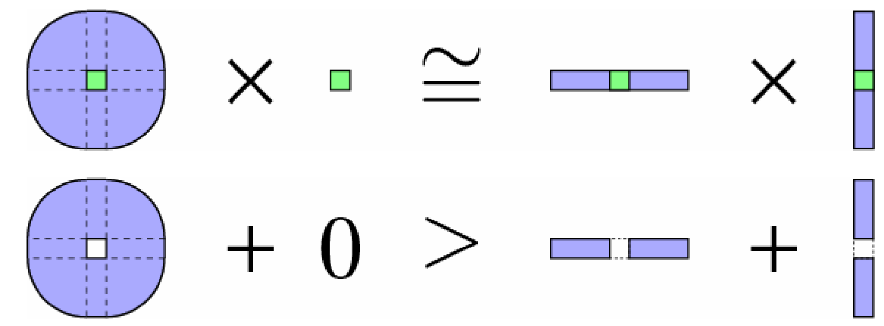

Traditional Method (2D) NT Method (2D)

The total amount of required computation per homebox is the same regardless of the strategy and is roughly proportional to the product of the volumes of the two "zones" of the strategy. The communication bandwidth is roughly proportional to the sum of the remote volume of these zones though. Traditional methods have unbalanced zone volumes. NT methods improve communication bandwidth by balancing the zone volumes while holding the product approximately constant. Even though most interactions are computed on a neutral homebox, the overall communications can be dramatically reduced in the limit of large numbers of homeboxes.

The Original NT Method

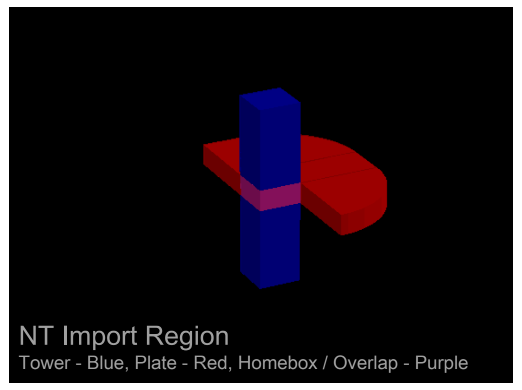

The import volume is:

$$V_i = 2\,R\,h_{xy}^{\;2} + 2\,R\,h_{xy}\,h_z + \pi\,R^2\,h_z\,/\,2$$

The volume of a homebox is:

$$V_h \;=\; h_{xy}^{\;2}\,h_z = V_s\,/\,p$$

The homebox aspect ratio ($h_{xy}\,/\,h_z$) is a free parameter. Optimizing to minimize import volume yields, in the limit of a large number of processors:

$$h_{xy} \rightarrow 2^{-1/2}\,(\pi\,R\,V_h)^{1/4}$$
$$h_z \rightarrow 2\,V_h\,(\pi\,R\,V_h)^{-1/2}$$
$$V_i \rightarrow 2\,\pi^{1/2}\,R^{3/2}\,V_h^{\;1/2} \sim O(\,R^{3/2}\,p^{-1/2}\,)$$

The x and y mesh coordinates of the homebox assigned a given interaction are that of the particle with the smaller x mesh coordinate. The z mesh coordinate is that of the other particle. The resulting import consists of an "outer tower" and "outer plate". Each homebox constructs a "tower" (set of particles in the outer tower and homebox) and "plate" (set of particles in the outer plate and homebox). Every tower particle is interacted with every plate particle (skipping various redundant interactions) with results exported to the appropriate homebox as necessary.

Snir's Hybrid Method and the SNT Method

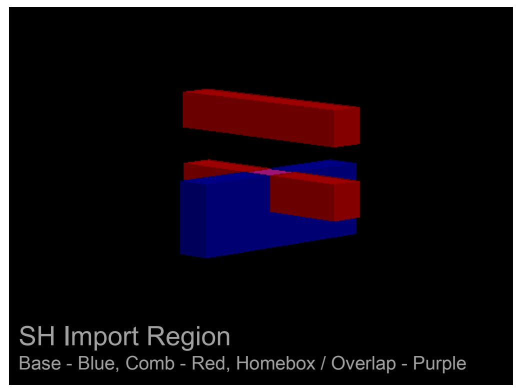

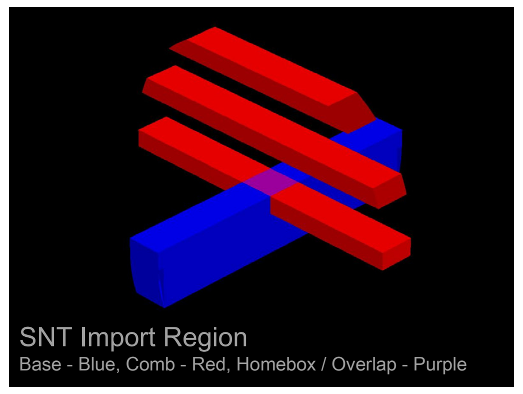

In Snir's hybrid (SH) method, "tower" and "plate" are replaced by "base" and "comb". The base is a connected region while the comb consists of a set of disjoint bars. The bar spacing is dual to the base thickness. The SH method assumed a cubic interaction cutoff and cubic homeboxes. It is possible to adapt the SH method for a spherical interaction cutoff and to optimize the homebox aspect ratio as was done in the NT method. This gives the SNT method.

Asymptotic Scaling of the HS, NT, SH and SNT Methods

Interaction Decomposition	**Asymptotic Bandwidth Scaling**
Atom	$O(\,1\,)$
Force	$O(\,p^{-1/2}\,)$
Traditional Spatial (HS)	$O(\,R^3\,)$
Shaw's Neutral Terriority (NT)	$O(\,R^{3/2}p^{-1/2}\,)$
Snir's Hybrid (SH)	$O(\,R^{3/2}\,p^{-1/2}\,)$ … ~13% greater asymptotic volume than NT
Snir's Hybrid with NT optimization (SNT)	$O(\,R^{3/2}\,p^{-1/2}\,)$ …~0.7% greater asymptotic volume than NT

p = number of processors, R = cutoff radius

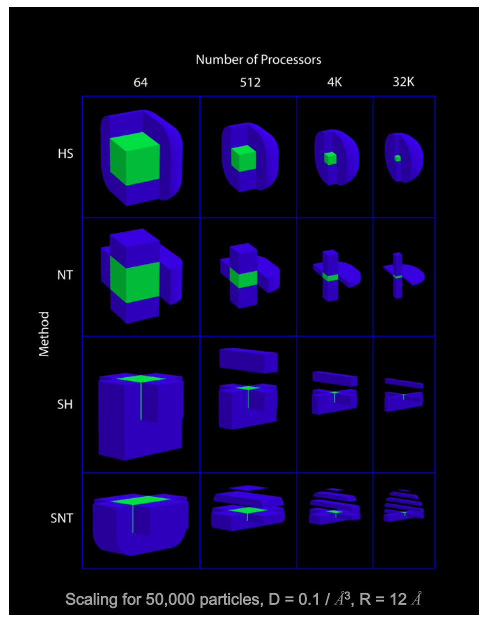

Generalized Decompositions

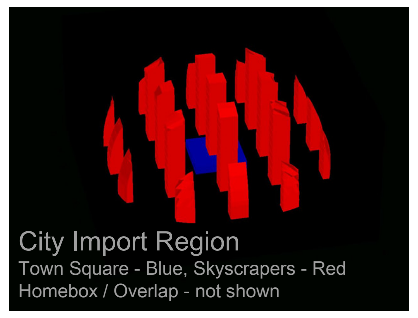

Convolution Criterion: Consider a method that assembles two sets ("zones") of complete homebox particle data. An interaction between two homeboxes is equivalent to the case where one homebox imported the other---like a "convolution". If the convolution of the zones covers half the shell of surrounding homeboxes (full if non-commutative), all necessary interactions will be computed at least once.

Rounding Criterion: Given a method that satisfies the convolution criterion, it can be improved by noting a point in a zone can be eliminated if it is further than R away from the closest point in the other zones with which the zone interacts. At left are two generalized NT methods that satisfy both criteria.

Assignment Rule: Alternatively, a method can be designed by giving a rule that uniquely determines the homebox responsible for an interaction and determining the necessary imports from the rule.

Multiple Zone Decompositions

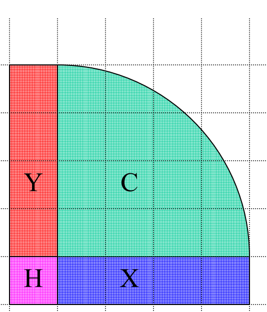

	H	X	Y	C
H	1*	2	3	4
X		0	5	0
Y			0	0
C				0

All methods considered so far have two zones that are interacted. There are many advantages to methods that use three or more zones though. To demonstrate such a method, consider the 2d commutative interaction assignment rule:

$$(i_{a\leftrightarrow b}, j_{a\leftrightarrow b}) = (\min(i_a, i_b), \min(j_a, j_b))$$

The zones corresponding to this rule do not overlap and are denoted H (in purple), X (in blue), Y (in red) and C (in green). This method is asymptotically inferior to the previous NT methods but it has equal or lower import volume than all known methods in the low parallelism limit. It cannot be trivially reduced to a method with fewer zones.

When there are multiple zones, the set of zone pairs to interact needs specification. It can be done with a sparse upper triangular matrix (the homebox-homebox interaction is a special case). The matrix gives the order in which to compute all necessary zone-zone interactions. Careful ordering can hide zone import behind computation.

The Multiple Zone NT Method

At right is shown a multiple zone formulation of the original NT method. H is the homebox, P is the outer plate, U is the upper tower and L is the lower tower. Application of the schedule to MPI-style cooperative communications is also shown. The schedule has eliminated redundancies of the original NT method. The schedule and rounding criterion also reveal an unexploited rounding opportunity. Because homebox plate-lower tower interactions in the original NT method were redundant, the lower tower can be additionally reduced in volume.

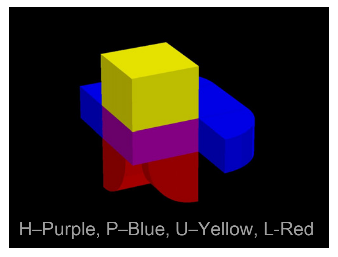

	H	P	U	L
H	1*	2	3	0
P		0	4	5
U			0	0
L				0

```
begin send of particles to P, U and L
begin receive of particles from P, U and L
begin receive of forces from P, U and L
do H-H interactions (special case)
end receive of particles from P
do H-P interactions
end receive of particles from U
do H-U interactions
do P-U interactions
begin send of forces to U
end receive of particles from L
do P-L interactions
begin send of forces to P and L
end receive of forces from P, U and L
```

A Lower Bound for Import Volume

Consider a commutative multiple zone method that has N_{zr} non-overlapping remote zones. The homebox treated as a separate, distinct zone and is the only self-interacting zone. If the method imports too little, it would be impossible to compute enough interactions. This gives a lower bound:

$$V_i \geq \frac{V_h}{\left(1-\frac{1}{N_{zr}}\right)}\left[\sqrt{\left(1-\frac{1}{N_{zr}}\right)\frac{V_{ir,remote}}{V_h}+1}-1\right] \text{ for } N_{zr} > 1$$

$$V_i \geq \left(\tfrac{1}{2}\right)V_{ir,remote} \text{ for } N_{zr} = 1$$

Introduce the parallelization parameter α_R and normalized homebox dimensions:

$$\alpha_R = R/V_h^{1/3}, \alpha_x = h_x/V_h^{1/3}, \alpha_y = h_y/V_h^{1/3}, \alpha_z = h_z/V_h^{1/3}$$

By considering various limits, $V_{ir,remote}$ can be approximated:

$$\frac{V_{ir,remote}}{V_h} \sim \frac{4\pi\alpha_R^3}{3} + \pi\alpha_R^2(\alpha_x+\alpha_y+\alpha_z) + 2\alpha_R(\alpha_x^{-1}+\alpha_y^{-1}+\alpha_z^{-1})$$

Minimizing over normalized homebox dimensions gives:

$$\frac{V_i}{V_h} \geq \left(1-\frac{1}{N_{zr}}\right)^{-1}\left\{\sqrt{\left(1-\frac{1}{N_{zr}}\right)\left[\frac{4\pi\alpha_R^3}{3}+3\pi\alpha_R^2+6\alpha_R\right]+1}-1\right\}$$

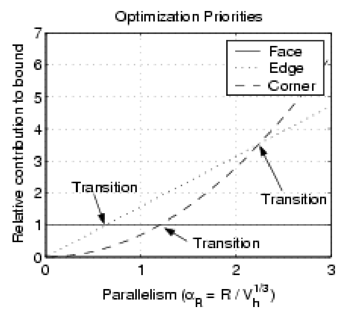

High parallelism bound (α_R>>*1*, N_{zr}>*1*, strict):

$$\frac{V_i}{V_h} \geq \left[\frac{4\pi\alpha_R^3}{3\left(1-\frac{1}{N_{zr}}\right)}\right]^{1/2} \rightarrow V_i \sim O\left(R^{3/2}p^{-1/2}\right)$$

Low parallelism bound (α_R<<*1*, approximate):

$$V_i \geq \left(\tfrac{1}{2}\right)V_{ir,remote}$$

Using more zones can mildly improve bandwidth

The Eighth Shell Method

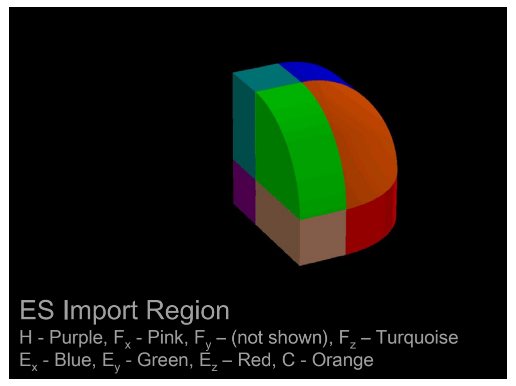

	H	F_X	F_Y	F_Z	E_X	E_Y	E_Z	C
H	1*	2	3	5	8	10	12	14
F_X		0	4	6	9	0	0	0
F_Y			0	7	0	11	0	0
F_Z				0	0	0	13	0
E_X					0	0	0	0
E_Y						0	0	0
E_Z							0	0
C								0

Above the import region of the eighth shell method is shown. It is the 3d generalization of the first multiple zone method given. It corresponds to the simple assignment rule:

$$(i_{a\leftrightarrow b}, j_{a\leftrightarrow b}, k_{a\leftrightarrow b}) = (\min(i_a, i_b), \min(j_a, j_b), \min(k_a, k_b))$$

This method is an optimal method in the low parallelism limit.

$$\frac{V_i}{V_h} = \frac{\pi\alpha_R^3}{6} + \frac{\pi\alpha_R^2}{4}(\alpha_x+\alpha_y+\alpha_z) + \alpha_R(\alpha_x^{-1}+\alpha_y^{-1}+\alpha_z^{-1})$$

The Split Half Shell Method

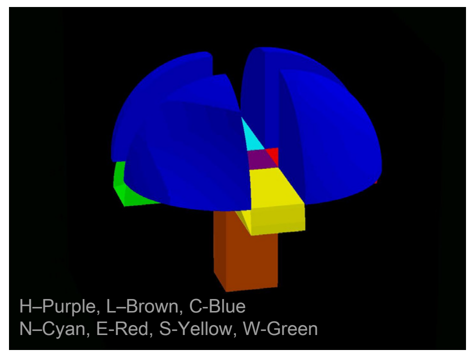

	H	L	N	E	W	S	C
H	1*	2	3	5	0	0	11
L		0	4	6	8	10	0
N			0	7	9	0	0
E				0	0	0	0
W					0	0	0
S						0	0
C							0

Above the import region of the split half shell method is shown. It has no straightforward assignment rule but can be proven by the convolution criterion. The method has "corner" and "face" imports but no "edge" imports. The optimization priorities suggest such methods are viable only over a narrow range of parallelization parameters (which is the case here).

$$\frac{V_i}{V_h} = \frac{2\pi\alpha_R^3}{3} + \alpha_R(2\alpha_x^{-1}+2\alpha_y^{-1}+\alpha_z^{-1})$$

The Foam Method

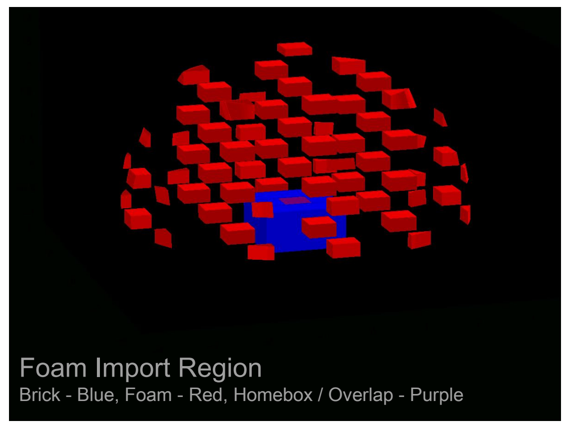

At left the import region of the foam method is shown. Though a multiple zone schedule is desirable in practice for redundancy elimination, it is omitted for brevity. Essentially, the foam method has two remote zones. A compact s_x x s_y x s_z homebox "brick" is interacted with a complementary disjoint "foam". This method achieves the lower bound import volume for a two zone method and is the best known method in the ultra high parallelism limit. The expressions below use a cubic brick of edge length s.

$$\frac{V_i}{V_h} \sim s^3 + \frac{2\pi R^3}{3s^3V_h} \sim \sqrt{\frac{8\pi R^3}{3V_h}} = \sqrt{\frac{8\pi\alpha_R^3}{3}}$$

Comparison of Methods

At right, the normalized import volume versus parallelization parameter is shown for the methods discussed here with cubic aspect ratio homeboxes. The Eighth Shell (ES) method is best for $\alpha_R < \sim 0.8$. The multiple zone variant of the original NT method (MZ-NT) is best for $\sim 0.8 < \alpha_R < \sim 2.7$. The multiple zone variant of Snir's hybrid method with NT optimizations (MZ-SNT) method is best for $\sim 2.7 < \alpha_R < \sim 14$. More exotic methods like foam become competitive for $\alpha_R > \sim 14$. Noting that:

$$\alpha_R = \frac{R}{V_h^{1/3}} = \frac{RD^{1/3}p^{1/3}}{N^{1/3}} \rightarrow p = \frac{\alpha_R^3 N}{R^3 D}$$

these ranges can be mapped to processor counts. For *50,000* particles, R = *12 Å* and D = *0.1 / Å³*:

ES	$p < \sim 150$
MZ-NT	$\sim 150 < p < \sim 5{,}700$
MZ-SNT	$\sim 5{,}700 < p < \sim 790{,}000$
Foam	$p > \sim 790{,}000$

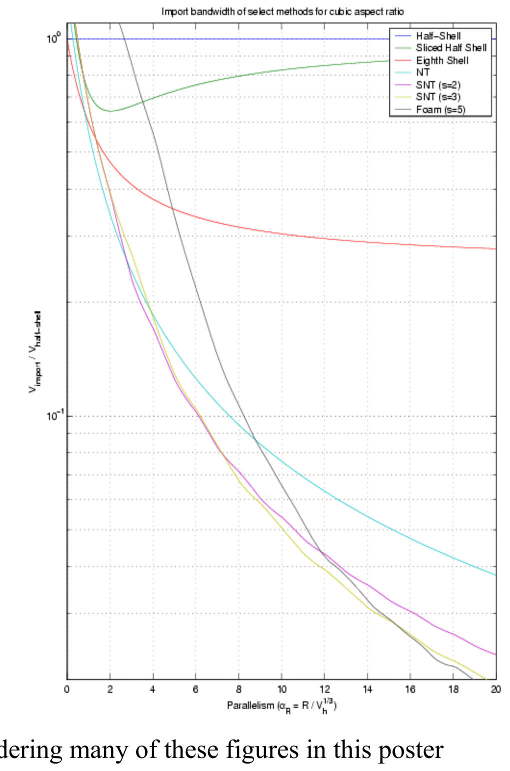

Special thanks to Christine McLeavey for rendering many of these figures in this poster